Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Model Uncertainty in Climate Change Economics

Loic Berger, Massimo Marinacci

Working Paper n. 616

This Version: November 13, 2017

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Model Uncertainty in Climate Change Economics*

Loïc Berger† Massimo Marinacci‡

November 13, 2017

Abstract

We review recent models of choices under uncertainty that have been proposed in the economic literature. The framework that we propose is general and may be applied in many different fields of environmental economics. To illustrate, we provide a simple application in the context of an optimal mitigation policy. Our objective is to offer guidance to policy makers who face uncertainty when designing climate policy.

1 Introduction

Uncertainty is pervasive. If this assertion is true for most decision problems, it is of particular importance when considering decisions with global, long-lasting and potentially irreversible consequences. The environmental challenge faced by humanity concerning global climate change illustrates particularly well the importance of considering uncertainty when making a decision. In this case indeed, decisions have to be made in the presence of uncertainty concerning both the science of climate and some basic socio-economic and technology drivers.

There is a growing awareness that the uncertainty encountered when dealing with problems like climate change goes well beyond the classical notion of "risk" typically used by economists. Put simply, the term *risk* refers to situations in which the probabilities of events' occurrence can be assumed to be known, while the notion of *uncertainty* is broader and refers to situations in which this may not be the case. Most decisions have indeed to be made in situations in which some events do not have an obvious, unanimously agreed upon, probability assignment. This might be the case because too little information is available,

*We are grateful to Valentina Bosetti, Laurent Drouet, Johannes Emmerling and Phoebe Koundouri for helpful comments and discussions. We acknowledge audiences at the annual conference of the European Association of Environmental and Resource Economists for their helpful comments. The research leading to these results has received funding from European Research Council (grants 336703 and 670337) and from the French Agence Nationale de la Recherche (ANR-17-CE03-0008-01).

†IESEG School of Management, LEM-CNRS 9221, Fondazione Eni Enrico Mattei (FEEM), Euro-Mediterranean Center on Climate Change (CMCC) and Università Bocconi. E-mail address: l.berger@ieseg.fr

‡Department of Decision Sciences and IGIER, Università Bocconi.

or because different predictions exist –resulting from different models or datasets, or from different experts' opinions.

Evaluation of climate policies is generally performed using models that do not make the distinction between risk and uncertainty, but actually reduce any kind of uncertainty to risk. The standard framework of expected utility theory developed by von Neumann and Morgenstern (1947) and Savage (1954) has been then used to explore rational decision making.

As the treatment of uncertainty has recently received a greater deal of attention in climate policy,[1] an increasing number of concerns have been raised about the use of standard techniques, originally developed to deal with risk, in problems involving uncertainty. For instance, the IPCC (2007) wrote:

> "In most instances, objective probabilities are difficult to estimate. Furthermore, a number of climate change impacts involve health, biodiversity, and future generations, and the value of changes in these assets is difficult to capture fully in estimates of economic costs and benefits. Where we cannot measure risks and consequences precisely, we cannot simply maximize net benefits mechanically. This does not mean that we should abandon the usefulness of cost-benefit analysis, but it should be used as an input, among others in climate change policy decisions. The literature on how to account for ambiguity in the total economic value is growing, even if there is no agreed standard."

As a result, recent calls have been made for using alternative tools and methods developed in other disciplines such as economics and statistics. One of them is made by Kunreuther at al. (2013):

> "The selection of climate policies should be an exercise in risk management reflecting the many relevant sources of uncertainty. Studies of climate change and its impacts rarely yield consensus on the distribution of exposure, vulnerability or possible outcomes. Hence policy analysis cannot effectively evaluate alternatives using standard approaches, such as expected utility theory and benefit-cost analysis. [...] For most issues relevant to policy choices, the solution is to use more robust approaches to risk management that do not require unambiguous probabilities. Risk management strategies designed to deal with the uncertainties that surround projections of climate change and their impacts can thus play an important role in supporting the development of sound policy options."

[1]See, for example, Pindyck (2007, 2013b), Heal and Miller (2014) and Convey and Wagner (2015) in recent issues of this journal.

In a recent *Science* article by Burke et al. (2016), twenty eight climate scientists outlined three areas where research progress on climate change economics was sorely needed. One of them consists in refining the estimates of the so-called *social cost of carbon* (SCC) to improve the way they are used in policy.[2] To achieve this objective, the authors highlighted different research directions, among which the treatment of uncertainty. They note:

> "The treatment of uncertainty in integrated assessment models needs improvement, with research needed on the computational challenges of explicitly including decision-making under uncertainty."

While the treatment of uncertainty has typically not received a particular attention in the environmental economic literature, the field is moving forward and several attempts have been made in the past few years to answer the aforementioned calls. These include, but are not limited to, Lange and Treich (2008) and Berger (2016) who provide comparative statics results of the role played by ambiguity in a simple two-period analytical model; Millner et al. (2013), Lemoine and Traeger (2016) who propose numerical climate-economic models under ambiguity aversion; Berger et al. (2017) who consider explicitly the presence of uncertainty concerning catastrophic climate events in both an analytical model and a numerical application; Athanassoglou and Xepapadeas (2012), Rudik (2016), Xepapadeas and Yannacopoulos (2017), who use robust control approach developed by Hansen and Sargent (2001, 2008) in either analytical control problems or in integrated assessment contexts; Drouet et al. (2015) who numerically disentangle model uncertainty and risks about mitigation costs, climate dynamics, and climate damages using the results of the most recent assessment of the Intergovernmental Panel on Climate Change (IPCC); Chambers and Melkonyan (2017) who compare three alternative decision criteria for climate change cost-benefit analysis in the presence of uncertainty; and Bradley et al. (2017), who deal with the uncertainty as presented by the IPCC by applying a recent model of Hill (2013) in which the confidence in the different models is not represented by a standard probability measure quantifying the uncertainty, but has rather a qualitative, ordinal, structure assessing DM's confidence in the probability judgements.

In this paper, we review recent models of choices under uncertainty that have been proposed in the economic literature and apply them to a simple climate change decision problem. While the framework we propose is general and may be applied in many different fields of environmental economics, we provide a simple illustrative application and example in the context of an optimal mitigation policy. Our objective is to offer guidance to policy

[2] The SCC is the damages caused by emitting carbon. In the words of Burke et al. (2016), the SCC estimates the "monetized change in social welfare over all future time from emitting one more tonne of carbon today, conditional on a specific trajectory of future global emissions and economic and demographic growth."

makers who face uncertainty when designing climate policies. In this regard, a recent related paper is Brock and Hansen (2017), who address, with a long term uncertainty perspective, some important climate policy issues by considering recent decision theoretic models.

Uncertainty can be decomposed into distinct layers: (i) *aleatory* or *physical uncertainty*, (ii) *model uncertainty* or *model ambiguity*, and (iii) *model misspecification*.[3]

Before elaborating on the key distinctions between these layers, let us have a closer look at the notion of "model uncertainty" since it may have different meanings depending on the field of analysis. In its colloquial sense, a "model" is generally viewed as a stylized –so, approximate but tractable– representation of a phenomenon of interest which a natural or social scientist wants to study. Models are used as tools that provide a logically consistent way to organize our thinking about the relationships among variables of interest, and help us understand the implications of those relationships (Maki, 2011, Pindyck, 2015, and Beck and Krueger, 2016). In environmental and climate change economics, a distinction is generally made between *scientific* models (climate and impact models), which inform us about the consequences of increased greenhouse gas (GHG) concentrations and emissions on the climate system as well as about the scale and nature of what might happen to lives and livelihoods, and *economic* models, which are used for cost-benefit analysis and policy assessments of alternative actions. A hybrid class of models, known as integrated assessment models (IAMs), combines the key elements of both economic and scientific models. They are generally used to calculate the SCC or to evaluate fiscal and abatement policies. These SCC estimates or evaluations (known as model runs) are then used directly by policy makers in cost-benefit analyses of climate change mitigation policies (Stern et al., 2016). There exist many models in all the different categories. Each model has its own advantages and limits, its own complexity, and its own key relationships and parameter values. Model uncertainty in the climate change literature is, therefore, usually associated with the fact that different models may provide different responses to the same external forcing (for instance, as a result of differences in physical and numerical formulations; cf. Deser, 2012).

The approach that we follow in this paper is, in part, different. We consider a general decision problems in which consequences depend on states of the environment that are viewed as realizations of an underlying economic or physical generative mechanism (Marinacci, 2015). A *model* is a probability distribution induced by such a mechanism. It describes states' variability by combining a structural component based on theoretical knowledge (say, economic or physical) and a random component coming, for example, from measurement errors or from minor omitted explanatory variables (cf. Koopmans, 1947, and Marschak, 1953). We assume that decision makers posit a collection of such mod-

[3]See Arrow (1951), Hansen (2014), Marinacci (2015), and Hansen and Marinacci (2016).

els. *Model uncertainty* therefore results from the uncertainty about the true underlying mechanism: within the posited collection, there is uncertainty about which model actually governs states' realizations. Once a model is specified, we still have the *aleatory uncertainty* about which specific state will actually obtain; this is the notion of *risk* typically considered in economics. Finally, we have a third layer of uncertainty, *model misspecification*, when the true model might not belong to the posited collection of models, so all posited models have an inherent approximate nature.[4]

An important instance of a similar approach in climate change economics concerns the estimations of climate sensitivity –the temperature change in response to increased atmospheric CO_2 concentration– presented in Millner et al. (2013) and Heal and Millner (2014). As these authors mention, climate sensitivity is an important metric for the study of climate change, yet it is difficult to estimate. Different complex scientific "models" attempt to predict its value but often do not agree one another. Each *scientific model* therefore delivers its own *probability model* for climate sensitivity (see Fig. 1 in Millner et al., 2013). Model uncertainty arises as uncertainty about the most correct scientific model, so about the true probability distribution for climate sensitivity.

The remainder of this paper is organized as follows. We first present a general decision problem under uncertainty framed in the context of climate change and discuss the different notions of uncertainty. We then present different models of choice that may be used to find the optimal climate policy, before applying them to a concrete example. We finish the paper with a discussion on the status of non-expected utility models in assessing climate policy.

2 The decision problem

2.1 Setup

An important challenge faced by environmental policy makers is to choose a mitigation strategy. Put simply, they have to decide how much GHG emissions should be allowed to avoid the climate system to reach damaging temperature levels. Reducing GHG emissions is costly, but it enables to limit the damages associated with global temperature increases. The cumulative level of GHG emissions an economy can reach over a given period of time (e.g., the twenty-first century) is called the carbon budget. It is a variable which is supposed to be directly under the control of the policy maker and is strictly related to global warming and climate targets (Meinshausen et al., 2009, Drouet et al., 2015). So, it is a decision (or control) variable –an *action* in the decision theory terminology– that represents a policy that the policy maker (or *decision maker*, hereafter DM) can perform.

[4]The notion of "true model" is, clearly, methodologically delicate. Pragmatically, here we consider the way such notion has been traditionally used in statistics.

At this point, it may useful to put some more structure on the decision problem under uncertainty. Formally, the problem that the DM faces is to choose an action a among a set A of possible alternatives, whose *consequences* c (within a consequence space C) depend on the realization of a *state of the environment* s (within a state space S) which is outside DM's control. The relationship between consequences, actions and states is described by a consequence function $\rho : A \times S \to C$, with

$$c = \rho(a, s). \tag{1}$$

In words, this function details the consequence c of action a if the state that obtains is s. DMs have a complete and transitive preference $\succsim$ over actions that describes how they rank the different alternative actions. In particular, we write $a \succsim b$ if the DM prefers action a to action b –i.e., either strictly prefers action a to action b, $a \succ b$, or is indifferent between the two, $a \sim b$.

The quintet $(A, S, C, \rho, \succsim)$ characterizes the decision problem under uncertainty. Before a decision is taken, *ex ante*, the DM knows the different elements of the quintet. After the decision, *ex post*, the DM observes the consequence $\rho(a, s)$ that obtained.[5] The objective of the DM is to select the action $\hat{a}$ that is *optimal* according to her preference in the sense that $\hat{a} \succsim a$ for all actions $a \in A$.

Preferences are often assumed to admit a numerical representation via a *decision criterion* $V : A \to \mathbb{R}$ such that

$$a \succsim b \Longleftrightarrow V(a) \geq V(b)$$

for all actions $a, b \in A$. The value $V(a)$ attained by an action a is often interpreted in welfare terms, so V is in turn viewed as a welfare criterion. However, here V is just a numerical representation of the underlying preference $\succsim$ that permits to formulate the decision problem as an optimization problem

$$\max_a V(a) \quad \text{sub } a \in A. \tag{2}$$

Coming back to the climate policy example, in principle the policy maker would like to set a global temperature increase to a level that maximizes a defined decision criterion. Global temperature, however, is not a decision variable under the control of the policy maker. In practice, what the policy maker controls is the level of emissions through an abatement policy that is put in place.

[5] Possibly not the state that obtained. In a dynamic setting, ex post observability becomes a key modelling issue (Battigalli et al., 2017).

2.2 Uncertainty

Decisions concerning the climate change phenomenon that our planet is undergoing are generally taken in situations of uncertainty. Think, for example, of a policy maker who has to choose the optimal emission pathway to be followed by an economy. It is reasonable to expect that the policy maker does not know, for example, how the climate system –in particular, global mean temperature– will respond to the targeted level of emissions, as well as how the socio-economic system will be affected by an increase in the global mean temperature. In that sense, the selection of the optimal action (level of emissions) is an exercise which is performed in a situation of uncertainty. The classical study of decision under uncertainty dates back to Savage (1954), while its modern study began with the behavioral paradoxes of Ellsberg (1961) and the theoretical analysis of Schmeidler (1989), with some key insights going back to Keynes (1921, 1936) and Knight (1921). Following the decomposition of climate change uncertainty according to its sources proposed by Heal and Millner (2014), the decision problem faced by our policy maker has to be performed in the presence of both a *scientific* and a *socio-economic* component. The next two examples will illustrate.

Scientific uncertainty A first source of uncertainty comes from the science of climate. While most scientists agree on the fact that climate change is a reality and that humans are primarily responsible for the unprecedented changes in global temperature that we have now been experiencing for several decades (Hansen et al., 2006; IPCC, 2013), the exact relationship between anthropogenic emissions of GHG into the atmosphere and climate change remains uncertain. Based on the available observations and on the current scientific state of knowledge, the scientific community has tried to construct precise climate models to predict and quantify the impact of human activity on global temperatures. Different metrics have been proposed to measure the global temperature response to increases in atmospheric emissions or concentrations. However, despite our knowledge about the physical laws governing the climate system,[6] a large degree of uncertainty still surrounds the estimates of these constructed climate metrics. For example, different scientific models typically provide different probabilistic assessments of the value of some key climate parameters (Meinshausen et al., 2009; IPCC, 2013; Millner et al., 2013). As a consequence, for instance we still do not know with certainty how much the global climate will exactly respond to changes in future atmospheric conditions, nor do we know the precise timing at which this change will take place.

The carbon-climate response (CCR) is an intuitive metric that has recently been proposed by Matthews et al. (2009) to provide policy-useful information about the allowable

[6]These laws for example enable us to narrow the scope of possible interactions between emissions and temperature increases.

level of emissions for a given temperature target. As illustrated in the next figure, this metric synthesizes the global temperature response to anthropogenic emissions. Formally,

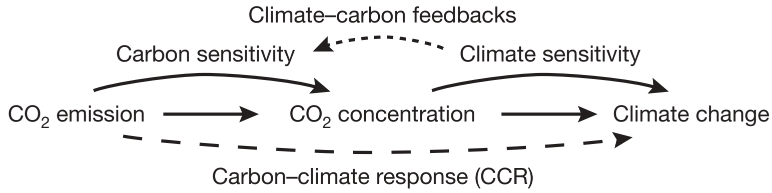

Figure 1: Schematic representation of the progression from CO_2 emissions to climate change.[7]

the CCR is defined as the ratio of temperature change to cumulative carbon emissions. It aggregates the well-known concept of climate sensitivity (the temperature change in response to increased atmospheric CO_2) and of 'carbon sensitivity' (the increase in atmospheric CO_2 concentrations resulting from CO_2 emissions as mediated by natural carbon sinks), while accounting for climate carbon cycle feedbacks. The CCR is claimed to be directly policy relevant, especially for climate change mitigation decisions. It combines the uncertainties associated with climate sensitivity, carbon sinks and climate-carbon feedbacks into a single metric.

Based on available historical data and observations, the CCR has been estimated by Matthews et al. (2009) to belong to the interval 1.0-2.1°C per trillion tones of carbon (TtC) for the period 1990-1999, with a best estimate of 1.5°C per TtC. The observational estimates of CCR are illustrated in the next figure. As can be observed, even when data

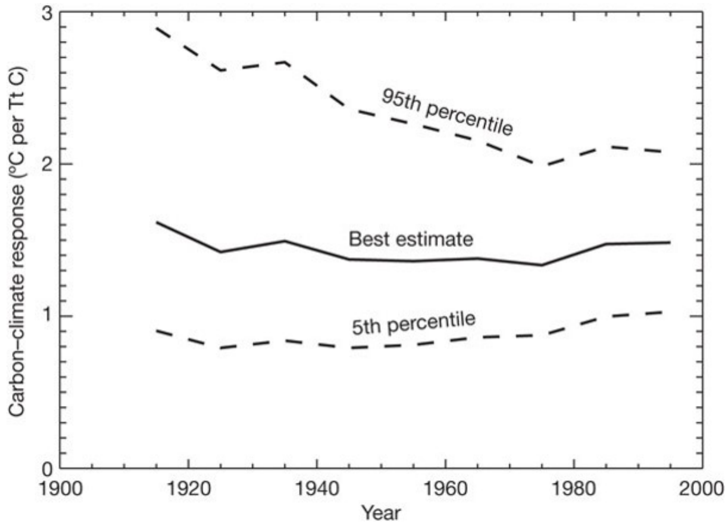

Figure 2: Observational estimates of CCR.[8]

about emissions and temperature changes are available, the exact relationship between

[7]Reprinted by permission from Macmillan Publishers Ltd: *Nature* 459: 829–832, copyright 2009.

the two cannot be established with certainty. We posit a stochastic linear relationship between emissions E and temperature increases T given by

$$T = \theta_T E + \varepsilon_T. \tag{3}$$

Here θ_T is a structural CCR parameter and ε_T is a random component that models the statistical performance of the underlying climate model due to measurement errors and to possible shocks.[9] It represents the unexplained variation caused by –possibly many– minor explanatory variables that we are "unable and unwilling to specify" (Marschak, 1953). The value of the CCR parameter has been shown to be remarkably constant within a given climate model, though it may varies across models due to differences in the way climate and carbon sensitivities are integrated. Figure 3 below reports the results of estimated CCR from eleven coupled climate-carbon cycle models participating in a model inter-comparison project (Friedlingstein et al., 2006). The figure presents the global temperature change as a function of cumulative carbon emissions. The relationship is remarkably linear for almost all the different models, thus justifying the linear form posited in (3). In this case, the CCR parameter θ_T for each model is nothing but the slope of the line that represents the intrinsic value of temperature change per unit of carbon emitted.

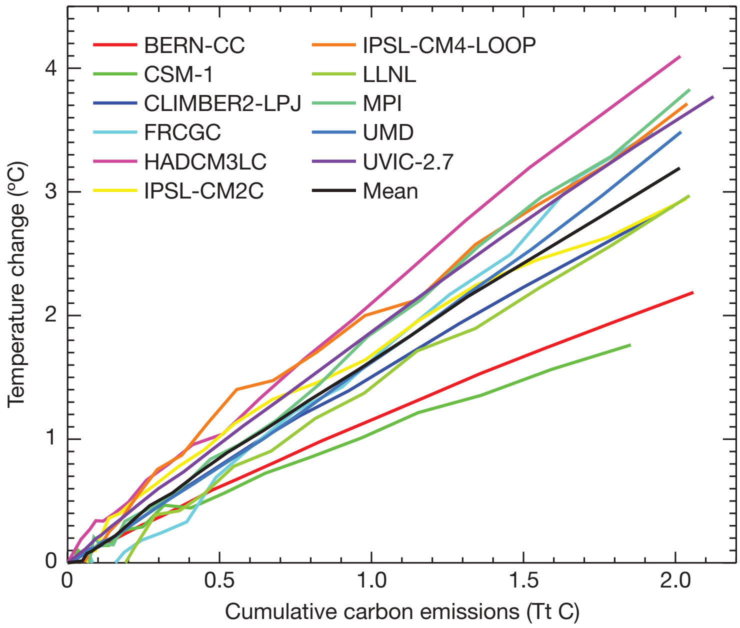

Figure 3: CCR estimated from different coupled climate-carbon simulation models.[10]

The value of the CCR for each model are presented in Table 1. As can be observed,

[9]A probability model quantifies the physical part of uncertainty, using "analogies with canonical random mechanisms that serve as benchmark" (Marinacci, 2015). So, we can regard random mechanisms as the thermometers of probability.

Model name	CCR (°C/TtC)	Model name	CCR (°C/TtC)
BERN-CC	1.1	IPSL-CM4-LOOP	1.9
CSM-1	1.0	LLNL	1.4
CLIMBER2-LPJ	1.5	MPI	1.9
FRCGC	1.7	UMD	1.6
HADCM3LC	2.1	UVIC-2.7d	1.9
IPSL-CM2C	1.7	**Mean**	**1.6**

Table 1: Values of the estimated CCR parameters from the eleven coupled climate-carbon models participating in the C4MIP project

model-based estimates of CCR range from 1.0 to 2.1°C per TtC.[11] Note that the uncertainty about the correct linear model (3), so about the CCR parameter, is epistemic. Within each such specification, there is still a layer of risk via the term ε_T.

Socio-economic uncertainty The second source of uncertainty faced by policy makers in the context of climate change concerns the relationship between global temperature increases and economic impacts. In an ideal world, physical and economic sciences should provide a theoretical underpinning for such relationship. In reality, the economic impact of global warming is complex and hard to predict (Pindyck, 2007). In the language of climate change economics, we have little information about the *damage function* d that represents this relationship between an increase in temperature T and the economic damage D or loss (usually expressed as fraction of GDP, see Pindyck 2013a, 2015). In other words, we do not have any economic or physical theory to help us assessing the “correct” functional form of this relationship. Moreover, since climate change mainly concerns events that we have never encountered so far, little data or empirical information can be used to assess both the degree of steepness of the damage function and the point where steepness begins. Traditionally, what integrated models of climate have been doing to deal with this problem is to use arbitrary functions to describe how GDP goes down when temperature increases. These functions, that rely on strong assumptions, have been subject to substantial criticism (Pindyck, 2013a; Howard et al., 2014; Howard, 2014), yet constitute the best approximation policy makers have at disposal. A typical damage function that has been used in this literature is the quadratic damage function

$$D = \theta_D T^2 + \varepsilon_D. \tag{4}$$

[11] The differences observed in the projection of the climate response to CO_2 emissions are due to different reasons such as different transient climate responses or carbon sensitivities as explained in the Supplementary Information of Matthews et al. (2009).

which is for example used in the DICE model of Nordhaus (1993); Nordhaus and Sztorc (2013).[12] The standard approach to calibrate this function has been to concentrate on the domain of relatively small increases in temperature, which is the only one for which we have some information at disposal. In the past few years, several studies have attempted to assess the impacts of global warming (or equivalently, the benefits from reducing GHG emissions).[13] These studies have used different methods –expert elicitation, enumeration, statistical methods, etc.– and have looked at different warming scenarios (usually within the range 1.0-3.0°C warming). The results of these studies, summarized in Table SM10-1 of IPCC (2014a), currently represent the best available information we have on the potential impacts resulting from climate change.

The lack of theoretical or empirical foundations concerning the damage function and its "correct" functional form does not matter too much when looking at relatively small temperature increases since there is a wide consensus that damages will be low at these levels. This is no longer the case for higher increases in temperature, which are associated with much stronger degrees of uncertainty. We have for example almost no idea of what damage to expect if temperature increases reach +5°C relative to preindustrial levels. Considering a temperature increase of $T = 5$ in (4) may therefore be misleading when analyzing climate policy given that calibration has been realized with data limited to small fluctuations in temperature taking place over relatively short periods of time (Pindyck, 2013a).

In a recent contribution, Drouet et al. (2015) summarize the information concerning total damages from global warming coming from the last IPCC report. These authors use twenty estimates of total economic effects of climate change to fit three different probabilistic damage functions. The results are presented in the next figure. Aggregating Drouet et al.'s (2015) specifications, the damage function can be represented by the following expression:

$$D = \theta_{1D}T + \theta_{2D}T^2 + \theta_{3D}T^6 + \theta_{4D}(\mathrm{e}^{-\theta_{5D}T^2} - 1) + \varepsilon_D, \tag{5}$$

for which three specifications of parameters are considered. When $\theta_{3D} = \theta_{4D} = 0$, the damage function has a quadratic form (first column of the figure) analogous to the one used in DICE and in most IAMs. When $\theta_{1D} = \theta_{2D} = \theta_{3D} = 0$ and $\theta_{4D} = 1$ (second column of the figure), we obtain a probabilistic version of the exponential damage function proposed by Weitzman (2009). This functional form is clearly steeper. It excludes the possibility of

[12] DICE stands for Dynamic Integrated Climate and Economy. To be precise, the damage function presented in equation (4) is the one used in the latest version of the DICE code (Nordhaus and Sztorc, 2013). It is a slight variant of the version of the quadratic form presented in the theoretical description of DICE, in which climate damages are bounded to 100% (i.e., climate change is assumed to only reduce current income, but may not destruct pre-existing assets). At low temperature increases, the two versions of the quadratic damage function are virtually identical.

[13] For an overview of these studies, see for example Pindyck (2013a); Heal and Millner (2014).

[14] Reprinted by permission from Macmillan Publishers.

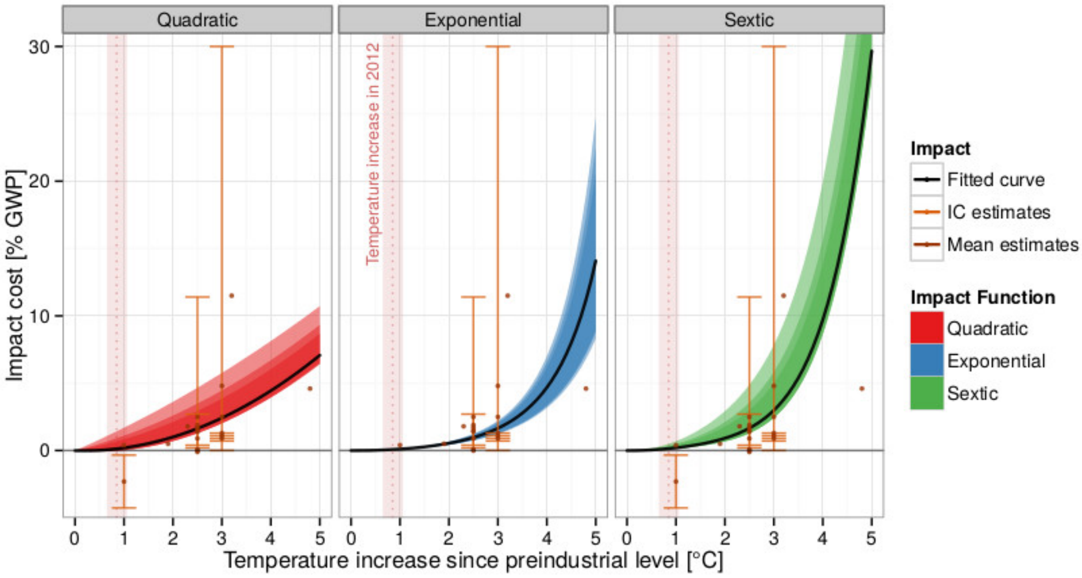

Figure 4: Probabilistic economic damage functions of temperature increase.[14]

potential benefits from climate change and allows for higher damages when temperature increases reach 4°C and 5°C. Finally, when $\theta_{1D} = \theta_{4D} = 0$ (third column of the figure), the damage function has the sextic form proposed by Weitzman (2012). According to this specification, high temperature increases are disastrous.

To illustrate the difference among the possible specifications of the damage function, consider Table 2 below. It presents the mean damages (and the 5$^{\text{th}}$-95$^{\text{th}}$ percentiles) associated with global warming expressed as percentages of world GDP, and obtained under the previous three different specifications of the damage function. The first row presents the results if temperature increases above preindustrial level reach +2°C. This is, indeed, the threshold that 195 countries have agreed to struggle for at the COP21. The second row presents the possible economic damages if the temperature increases reach +3°C. This level of warming roughly corresponds to the median 2100 temperature increase projection if "nationally determined contributions" (NDCs; i.e., climate pledges that each country made to tackle the problem of climate change) are implemented as planned (Bosetti et al., 2017).[15] Finally, the last two rows concern the more extreme temperature increases of +4°C and +5°C. These levels of warming roughly determine the bounds of the temperature changes we could expect under a business-as-usual situation (i.e., if no additional effort are made to constrain emissions; IPCC, 2014b).

These three damage functions do not have a clear theoretical underpinning, they just fit the best data currently available on potential losses using different specifications. The

[15]Models used for projections of future temperature increases are the ones whose results on transient climate response are reported in the IPCC fifth assessment report. The hypothesis that current NDCs are projected beyond 2030 is here made for these projections. See Bosetti et al. (2017) for more details.

Temperature increase since preindustrial level	Type of damage function		
	Quadratic	Exponential	Sextic
2° C	1.01 [0.74; 2.09]	0.54 [0.44; 0.63]	0.49 [0.37; 0.90]
3° C	2.42 [2.04; 3.81]	1.66 [1.28; 2.02]	1.93 [1.72; 2.73]
4° C	4.45 [3.96; 6.00]	4.68 [3.34; 6.14]	8.5 [8.04; 10.49]
5° C	7.07 [6.49; 8.63]	14.09 [8.93; 20.56]	30.54 [29.40; 36.14]

Notes: Mean damages, 5th-95th confidence interval in brackets

Table 2: Economic damages from climate change (in % of world GDP)

random component ε_D accounts for omitted variables and measurement errors. There is, therefore, a high degree of structural uncertainty regarding the "correctness" of the functional form representing this relationship. This uncertainty is of the epistemic nature: the policy maker does not know which is the most accurate model to describe the relationship between global warming and GDP among the three potential models proposed by economists. The probability that may be attached to each model is therefore a representation of the policy maker's degree of belief. Yet, a layer of risk is also present within each model via the term ε_D.

2.3 Decision under uncertainty

Let us now go back to the decision problem faced by our policy maker, who has to choose the action representing the level of emissions –i.e., $a = E$– knowing that it will affect global temperatures via the carbon-climate equation (3) which, in turn, will affect the economic output via a damage function.

To simplify the computation, let us assume that damages are represented by the quadratic equation (4). As we argued in the previous section, the relevant scientific and socio-economic relationships may be summarized by the following nonlinear system:

$$\begin{cases} T = \theta_T E + \varepsilon_T \\ D = \theta_D T^2 + \varepsilon_D \end{cases} \tag{6}$$

States have, in general, both random and structural components, so they have the form

$$s = (\varepsilon, \theta).$$

In our decision problem, the vector $\varepsilon = (\varepsilon_T, \varepsilon_D)$ represents the pair of random shocks affecting the climate and economic systems, while the vector $\theta = (\theta_T, \theta_D)$ specifies the structural coefficients, parametrizing via θ_T a model climate system and via θ_D a model economy. By substitution from the system (6), the damage function has thus the form

$$d(E, \varepsilon, \theta) = \theta_D \theta_T^2 E^2 + \theta_D \varepsilon_T^2 + 2\theta_D E \varepsilon_T + \varepsilon_D.$$

For each policy E, the consequence function

$$\rho(E, \varepsilon, \theta) = -d(E, \varepsilon, \theta) - c(E)$$

specifies the overall monetary outcome in terms of some economic variable of interest –e.g., consumption or GDP– given the random components and the monetary cost $c(E)$ of the policy itself.

We suppose that, following the ex ante scientific and socio-economic information, the policy maker is able to posit a set of potential models M describing the likelihoods of the different states. This set of models is taken as a datum of the decision problem: the policy maker behaves as if she knows that states are generated by a probability model m that belongs to the collection M. We thus abstract from model misspecification, which of course would magnify the issues that we will discuss.

The positive scalar $m(\varepsilon, \theta)$ gives the joint probability of shock ε and parameter θ. It is natural to adopt the factorization $m = q \times \delta_\theta$, that is,

$$m(\varepsilon, \theta') = \begin{cases} q(\varepsilon) & \text{if } \theta' = \theta \\ 0 & \text{else} \end{cases} \tag{7}$$

where $q(\varepsilon)$ is the probability of ε and δ_θ is the probability distribution concentrated on θ.[16] Each model corresponds to a shock distribution q of the vector ε and to a model climate system/economy θ. In this factorization,[17] two kinds of model uncertainties emerge. First, there is uncertainty about the economic and physical theories that underpin the models: different θs correspond to different such theories. Second, there is uncertainty about the statistical performance of such theories, due to shocks and to measurement errors: different q correspond to different such performances. We may call, respectively, *theoretical model uncertainty* and *stochastic model uncertainty* the two types of uncertainty. The former is more fundamental than the latter because it reflects the scientific (natural and social) views of policy makers. For this reason, we assume that the shocks' distribution q is known and common across models. If so, the only epistemic uncertainty that remains is about the structural component θ. Different models thus correspond to different specifications

[16] That is, $\delta_\theta(\theta) = 1$ and $\delta_\theta(\theta') = 0$ if $\theta' \neq \theta$.

[17] In a monetary policy context, a similar factorization was assumed in Battigalli et al. (2016).

of the structural component, so we can index the models via the structural parameters

$$m_\theta = q \times \delta_\theta$$

and write $M = \{m_\theta\}_{\theta \in \Theta}$, where m_θ parametrizes model (7).[18] Given a structural parameter θ, aleatory uncertainty is quantified by the known distribution q. To address the epistemic uncertainty about θ, the DM may have a subjective prior probability distribution μ that quantifies her belief –based on her personal information– about the true structural parameter. So, $\mu(\theta)$ is the DM's subjective belief that θ is the true parameter, that is, that m_θ is the true model.

Now that all the elements of the decision problem under uncertainty have been introduced, let us turn to the way they can be combined to make the best possible decision. For this purpose, we describe different decision criteria developed in economic theory that may be used for problems of decision under uncertainty.

3 Classical subjective expected utility

We start with the description of the decision criterion that, for several decades, has been viewed as the standard way to consider rational decision making under uncertainty. This criterion, which dates back to the seminal works of von Neumann and Morgenstern (1947), Wald (1950), Savage (1954), and Marschak and Radner (1972), has recently been revisited by Cerreia-Vioglio et al. (2013) to accommodate explicitly the presence of model uncertainty.

Let us consider the decision problem $(A, S, C, \rho, \succsim)$ defined in Section 2. Assume that a von Neumann-Morgernstern utility function $u : C \to \mathbb{R}$ translates economic consequences, measured in monetary terms, into utility levels. As is well-known, this function captures risk attitudes (i.e., attitudes towards aleatory uncertainty). For each action a and each model m_θ characterizing a combination of climate/economy environment, it is therefore possible to compute the *expected reward* (or *payoff*) associated with a given action:

$$R(a,\theta) = \sum_{s \in S} u\left(\rho(a,s)\right) m_\theta(s). \tag{8}$$

For instance, under risk neutrality we have

$$R\left(E,\theta\right) = -\theta_D \theta_T^2 E^2 - \theta_D - c(E)$$

[18] If we drop the assumption that q is known, we have $m_{\chi,\theta} = q_\chi \times \delta_\theta$ where χ parametrizes the possible distributions q. In this case, the belief is over two parameters, that is, $\mu\left(\chi,\theta\right)$. An intermediate case is when the distribution q is assumed to be know but may vary across models. In this case, we have $m_\theta = q_\theta \times \delta_\theta$ so the belief is still on a single parameter, that is, it has the form $\mu\left(\theta\right)$.

provided the random components have zero mean and unit variance.

Given that different models exist and that the DM does not know which is the correct one, she considers the expected payoff of each possible model and aggregates them out by performing a weighted average according to the relative weights that she associates with each of them (i.e., her prior probability μ). The *classical subjective expected utility* decision criterion is:

$$V_{eu}(a) = \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta). \tag{9}$$

The *optimal policy* therefore consists in choosing the action $\hat{a}$ that maximizes this criterion.[19] Formally, this amounts to solving the optimization problem (2), which here takes the form

$$\max_{a} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta) \quad \text{sub } a \in A \tag{10}$$

Optimal actions therefore depend on DM's preferences via the utility function u and the prior probability μ. If the prior distribution is uniform, criterion V_{eu} consists in an average of the expected rewards $R(a, m_\theta)$ where all the models are equally weighted. The optimal policy maximizes such average expected payoff.

Criterion (9) is a Bayesian two-stage criterion that describes both layers of uncertainty, risk and model uncertainty, via standard probability measures. It is possible to write it à la Savage (1954) as a single stage criterion:

$$V_{eu}(a) = \sum_{s \in S} u(\rho(a, s))\, \bar{m}(s), \tag{11}$$

where $\bar{m}(s) = \sum_{\theta \in \Theta} m_\theta(s)\, \mu(\theta)$ is the so-called *predictive distribution* on states. Under a uniform prior, in our example $\bar{m}$ would correspond to the model that features the mean CCR and the mean damage. In words, the equality between expressions (9) and (11) tells us the following: considering a collection of models M and aggregating them is the same as considering a unique average model. This is made possible only because the same attitude is considered towards both risk and model uncertainty (see Marinacci, 2015). To see this, note that criterion (9) may be rewritten as

$$\begin{aligned} V_{eu}(a) &= \sum_{\theta \in \Theta} \left(u \circ u^{-1}\right)(R(a, \theta))\, \mu(\theta) \\ &= \sum_{\theta \in \Theta} u\left(u^{-1}(R(a, \theta))\right) \mu(\theta) \end{aligned}$$

[19] To ease matters, we restrict our attention to finite state and model spaces. Integrals with respect to probability density functions would arise without such assumption.

in which it is made clear that the outer u represents attitude towards model uncertainty, while the inner one –entering the monetary certainty equivalent $u^{-1}(R(a,\theta))$– represents risk attitude. In that sense, criterion (9) overshadows the DM's reaction to the variability that may exists across models. Considering equally the different CCRs or a single CCR= 1.6 on which everyone would agree, for example leads exactly to the same optimal emission policy. Indeed, these very different scenarios reduce to the same predictive model $\bar{m}$. Criterion (9) therefore presupposes that the policy maker has the same attitude towards aleatory and epistemic uncertainty.

Before relaxing, in Section 6, the assumption of equal treatment between the different layers of uncertainty, let us consider two important special cases of criterion (9). First, suppose that the DM considers (wrongly, possibly) only a parameter to be the correct one, so there is a 100% weight on it; formally, $\mu(\theta) = 1$. The two-stage criterion (9) then reduces to $V_{eu}(a) = R(a, \theta)$. In this case, model uncertainty is still part of the decision problem, but the DM is dogmatic about a specific model being the correct one and therefore does not take into account any other model. Second, suppose that there is only one parameter in the collection Θ, for instance there is no scientific uncertainty about the value of the CCR parameter nor economic uncertainty about the correct damage function. In this case, the DM knows that θ is the correct parameter. Epistemic uncertainty is not anymore present in the decision problem, which is a decision problem under risk –as represented by model m_θ. In this case, criterion (9) reduces to $V_{eu}(a) = R(a,\theta)$ interpreted as a von Neumann-Morgenstern risk criterion. This is typically what is implied by the *rational expectations hypothesis* often adopted in economics, which assumes that DMs know the correct model.

We now move beyond the classical subjective expected utility criterion (9) and discuss alternative decision criteria under uncertainty.[20]

4 Unanimity preferences

One way to deal with uncertainty is to allow for preferences to be incomplete. Because of the lack of knowledge concerning both the science of climate and the impact of climate change on the economy, the policy maker might not be able to rank some pairs of alternative actions. If this is the case, the preference $\succsim$ is no longer complete (as so far assumed) but *incomplete*. Assume, following the classic analysis of Bewley (2002), that the DM knows her tastes and is able to rank any pair of consequences, so there is a well defined utility function u. Imagine, however, that the DM is unable to rank some pairs of actions because of insufficient information about them. Because of its incompleteness, the preference $\succsim$ cannot be represented via a numerical decision criterion V, but via a non-numerical unanimity rule:

[20]Note that these decision criteria have axiomatic behavioral foundations that clarify their nature. We refer interested readers to Gilboa and Marinacci (2013).

$$a \succeq a' \iff R(a,\theta) \geq R(a',\theta) \qquad \forall \theta \in \Theta. \tag{12}$$

In words, action a is preferred to another action a' if and only if, according to all the probability models m_θ, the expected reward associated with action a is higher than that associated with action a'.[21] In our emission example, this would be the case if and only if policy a is better than policy a' according to all the different climate/economy models.

A unanimity criterion is often unable to specify what the DM should do. This is the case for example if a policy (say, a low level of emission policy) performs better than another policy (say, a high level of emissions policy) according to some models, but performs worse according to other models. If, nevertheless, a decision has to be made, one needs to "complete" the criterion when it remains silent. A possibility is to adopt a default decision rule that relies on a *status quo action* that remains the default action until it is replaced by an alternative action that is unanimously better. In climate change economics, this status quo situation has typically consisted in the "wait-and-see" policy, and uncertainty has long been seen as an excuse for inaction in climate policy. Other approaches suggest to complete preferences with one of the criteria that are presented below (Gilboa et al., 2010, Cerreia-Vioglio, 2016), so to take care of the burden of choice in a less ad hoc manner than status quo.

In sum, the unanimous criterion (12) may turn out to be useless in situations where a choice has to be made. In the next sections, we will present alternative criteria that preserve completeness, so the numerical nature of the decision criterion. In particular, we will relax the assumption of the classical subjective expected utility theory according to which risk and model uncertainty attitudes are both captured by the same function u. In principle, there is indeed no reason to expect these two attitudes to be equal. A policy maker might well be more prone to be confronted with the risk due to the intrinsic randomness of some events than to be confronted to the model uncertainty due to a lack of, say scientific, knowledge. Recent experimental evidence on both students and real-life policy makers shows that this is indeed the case (Berger and Bosetti, 2017). A policy maker fulfilling this condition is said to be more averse to model uncertainty than to risk, and consequently exhibits uncertainty (or ambiguity) aversion. This latter behavioral characteristic, first highlighted by Ellsberg (1961), has been shown to robustly describe the behavior of individuals in situations of uncertainty.

[21] The unanimity criterion (12) is based on the general form of Bewley's model studied by Gilboa et al. (2010), but it is conceptually different in that here Θ is posited.

5 Classical analysis

The maxmin criterion of Wald (1950) is a first decision criterion that considers attitudes towards risk and model uncertainty differently. This criterion is extremely cautious because it makes the DM to consider only the model giving her the lowest expected payoff. In our examples, this means that only the "worst" out of the possible climate/economy models is considered when choosing the optimal climate policy. Prior probabilities do not play any role here, so we are in a classical statistics setting. Formally, the *maxmin decision criterion* is:

$$V_{mxm}(a) = \min_{\theta \in \Theta} R(a, \theta). \tag{13}$$

Choosing the optimal policy under this criterion corresponds to finding the value of a that maximizes the minimal expected payoff obtained over the set of possible probability models (this is why this criterion is called "maxmin"). This criterion has, for example, recently been used by Rezai and van der Ploeg (2017) to study the impact of scientific uncertainty regarding the "correct" climate models to be used in the context of integrated assessment models.

6 Uncertainty averse preferences

6.1 Bayesian analysis

Another way to distinguish attitudes to model uncertainty and to risk is to adapt the smooth ambiguity model developed by Klibanoff et al. (2005). In that case, the *smooth decision criterion* to be maximized is:

$$V_{smt}(a) = \sum_{\theta \in \Theta} \phi \left(R(a, \theta) \right) \mu(\theta), \tag{14}$$

where $\phi \equiv v \circ u^{-1}$ represents the attitude towards uncertainty that results from the combination of attitudes towards model uncertainty v and towards risk u. Concavity of ϕ reflects uncertainty aversion that, in this setup, amounts to a stronger aversion to model uncertainty than to risk –i.e., v is more concave than u (cf. Marinacci, 2015).

Like (9), also the smooth ambiguity criterion (14) is a two-stage Bayesian criterion in which both layers of uncertainty are described by standard probability measures. It may be written as $\sum_{\theta \in \Theta} v \left(u^{-1} \left(R(a, \theta) \right) \right) \mu(\theta)$ and interpreted as follows. In the first stage, the DM evaluates the expected payoff of policy a per each possible model m_θ and expresses it in monetary terms through a certainty equivalent $c_\theta \equiv u^{-1}(R(a, m_\theta))$. The certainty equivalents c_θ represent the amount of the economic variable of interest –e.g., GDP or consumption– which would make the DM indifferent between getting such amount for sure and facing the risk that model θ involves. A certainty equivalent c_θ may be computed per each model. It depends on risk attitude via the function u: the more risk averse the DM is, the lower c_θ is. In the second stage the DM addresses model uncertainty, the decision theoretic scope of which is described by the collection $\{c_\theta : \theta \in \Theta\}$ of certainty equivalents. The DM summarizes the welfare impact of policy a by evaluating an overall expected payoff $\sum_{\theta\in\Theta} v(c_\theta)\mu(\theta)$ across the certainty equivalents by using her attitude towards model uncertainty v and her prior belief μ. This is exactly what represents the two-stage decision criterion (14).

As before, if the prior distribution is uniform, the certainty equivalents are given the same weight in computing the overall expected welfare. If model uncertainty in the second stage is evaluated using risk attitude u, so $u = v$, we are back to the classical subjective expected utility criterion (9), which corresponds to a situation of ambiguity neutrality (function ϕ is linear in this case, thus the overall expected welfare is the average of the expected rewards). Interestingly, the maxmin criterion (13) is a limit case of the smooth decision criterion (14) when model uncertainty tends to infinity (see Klibanoff et al., 2005). For instance, if $\phi_\lambda(x) = -e^{-\lambda x}$ we have[22]

$$\lim_{\lambda\to+\infty} \phi_\lambda^{-1}\left(\sum_{\theta\in\Theta} \phi_\lambda\left(R\left(a,\theta\right)\right)\mu(\theta)\right) = \min_{\theta\in \mathrm{supp}\mu} R\left(a,\theta\right),$$

which reduces to (13) when μ has full support, i.e., $\mathrm{supp}\mu = \Theta$. As here uncertainty aversion results from higher aversion to model uncertainty than to risk, it should be clear that the maxmin criterion corresponds to an extreme aversion to model uncertainty relative to risk.

In a recent contribution, Berger et al. (2017) explicitly made the distinction between attitudes towards aleatory and epistemic uncertainty while using the smooth criterion to study the impact of scientific uncertainty regarding the possibility of a particular climate catastrophe on the optimal level of GHG emissions. Another example of application of this criterion in climate change economics may be found in the work of Millner et al. (2013).

6.2 Non Bayesian analysis

We already performed non Bayesian analysis when presenting Wald's maxmin criterion, where priors play no role. Yet, a different departure from the Bayesian framework originates in the work of Gilboa and Schmeidler (1989).

Multiple priors The multiple priors approach relaxes the assumption that the DM's information about model uncertainty is quantified through a single probability distribution

[22]See Klibanoff et al. (2005), who interpret the coefficient λ can be interpreted in terms of absolute coefficient of ambiguity aversion.

μ. Instead, it allows for the possibility that it is quantified by a set C of them because the DM does not have sufficient information to specify a single prior over the different models. The *multiple priors decision criterion* is

$$V_{mp}(a) = \min_{\mu \in C} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta). \tag{15}$$

Contrary to Wald's extreme criterion –with which it is sometimes confused– the multiple priors criterion of Gilboa and Schmeidler (1989) considers the least favorable among all the classical subjective expected utilities determined by each prior μ in C. In our climate policy example, a particular prior distribution may be the uniform that gives equal weight, $\mu(\theta) = 1/|M|$, to all the possible models, while another prior may for example not consider some values of the CCR as plausible (in which case, some $\mu(\theta)$ have a value 0). The classical subjective expected utilities is then computed for each prior distribution, with the optimal policy being the one that maximizes the expected payoff obtained with the "worst prior".

Criterion (15) has also often been called the maxmin criterion, yet it is less extreme than it may appear at a first glance. The set C of possible priors incorporates both the attitude towards uncertainty and an information component: a smaller set C may reflect both better information and/or less uncertainty aversion. In any case, a more general and less extreme α-version of this model has been axiomatized by Ghirardato et al. (2004) in which both the "max" and the "min" appear with weights α and $1-\alpha$. This more general form

$$V_{\alpha\text{-}mp}(a) = \alpha \min_{\mu \in C} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta) + (1-\alpha) \max_{\mu \in C} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta)$$

may accommodate milder, even positive, attitudes toward uncertainty.

Two criteria that we have already encountered are special cases of the multiple priors model. First, the classical subjective expected utility criterion (9) is recovered when the set C is singleton (i.e., it contains only one element). Second, we return to Wald's maxmin criterion (13) when the set C is maximal in that it consists of the set $\Delta(\Theta)$ all possible prior probabilities. Indeed, we have

$$\min_{\mu \in \Delta(\Theta)} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta) = \min_{\theta \in \Theta} R(a, \theta).$$

So, Wald's maxmin can be interpreted as the extreme case of maximal "prior uncertainty".

Robustness Another criterion, known as the *variational decision criterion*, has been axiomatized by Maccheroni et al. (2006). It is written as:

$$V_{vr}(a) = \min_{\mu \in \Delta(\Theta)} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta) + c(\mu)$$

Under this criterion, priors are weighted by a convex function c. Importantly, if c is strictly convex, the criterion V_{vr} becomes differentiable. This criterion has a penalization form familiar from robust control theory. In particular, if c is dichotomous, being 0 if μ belongs to some set C and $+\infty$ otherwise, we are back to the multiple priors criterion (15). In contrast, if c has the relative entropy form $\lambda^{-1}R(\mu \parallel \nu)$ with respect to a reference prior ν and a coefficient $\lambda > 0$, we have the *multiplier decision criterion*

$$V_{rb}(a) = \min_{\mu \in \Delta(\Theta)} \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta) + \frac{1}{\lambda} R(\mu \parallel \nu)$$

of Hansen and Sargent (2001, 2008). Because of a convex analysis equality (Dupuis and Ellis, 1997, p. 27), the multiplier decision criterion can be equivalently written in the smooth ambiguity form

$$V_{rb}(a) = \phi_\lambda^{-1}\left(\sum_{\theta \in \Theta} \phi_\lambda\left(R(a, \theta)\right) \mu(\theta)\right) = -\frac{1}{\lambda} \log \sum_{\theta \in \Theta} e^{-\lambda R(a, m_\theta)} \mu(\theta)$$

Indeed, as noted by Hansen and Sargent (2007) and Cerreia-Vioglio et al. (2011), the multiplier decision criterion is, essentially, the intersection of smooth ambiguity averse and variational representations.[23]

Examples of applications in climate change economics of the multiplier decision criterion may, for example, be found in the works of Athanassoglou and Xepapadeas (2012), Rudik (2016), and Xepapadeas and Yannacopoulos (2017).

6.3 Other approaches

The approaches discussed so far have a normative motivation. They assume that DMs have to cope with uncertainty without expecting to reduce everything to risk, a pretension that tacitly presumes a much better information than they typically have. Making decision under a fictitious, even delusional, state of information seems hardly a rational way to proceed. That said, other approaches have been proposed with a descriptive motivation. These include, for example, prospect theory (see Wakker, 2010). However, their descriptive motivation make them less relevant for the climate policy problem that we consider.

Finally, note that another criterion known as *minmax regret* and due to Savage (1951) is also sometimes used in the environmental literature. Because it violates the independence of irrelevant alternatives, a basic rationality tenet, we however do not discuss this criterion here and refer the interested reader to a discussion provided in Marinacci (2015).

[23] These "robust" criteria can be set also in a classical setting, without priors. For instance, the variational decision criterion becomes $V(a) = \min_{\theta \in \Theta} R(a, \theta) + c(\theta)$ in a classical setting.

6.4 Summing up

To sum up, we can represent the previous numerical decision criteria in the following diagram.

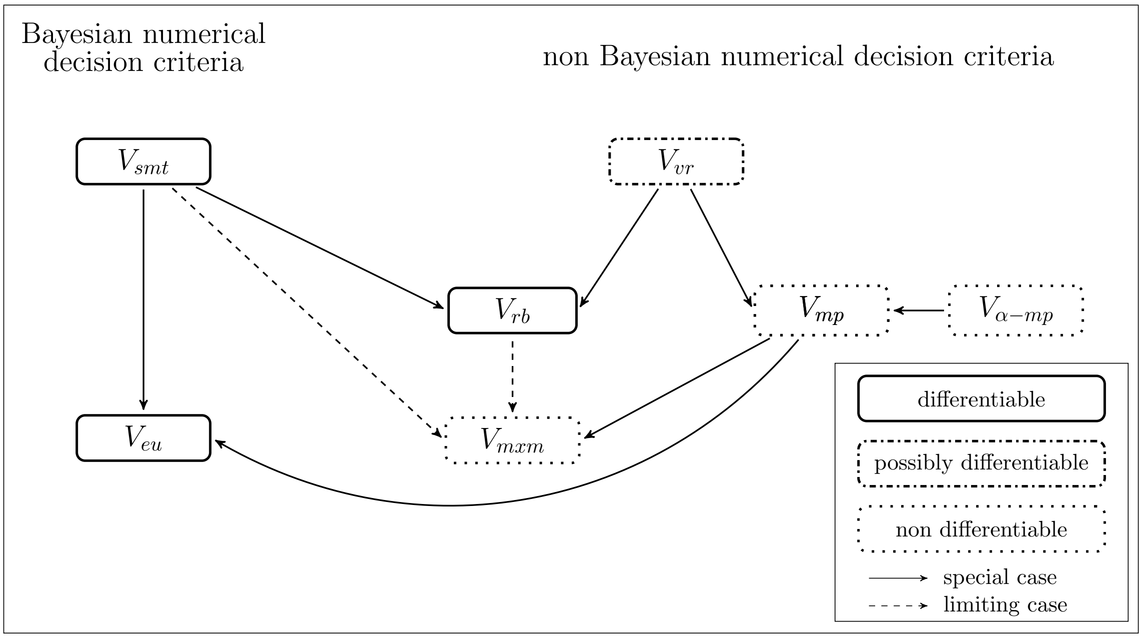

Figure 5: Representation of alternative decision criteria

7 Application

To illustrate the differences in terms of optimal climate policy among the distinct criteria that we presented, let us go back to our example of a policy maker looking for the optimal mitigation policy to put in place. For the sake of simplicity, let us assume that her objective is to choose the level of GHG emissions maximizing the net level of output of the economy. This level simply corresponds to the level of output net of the damages due to climate change and of the costs necessary to reduce emissions (the so-called abatement costs). Because of the presence of scientific and socio-economic uncertainties, the net level of output is itself uncertain.[24] Thanks to the best available scientific and economic information at her disposal, the policy maker knows that thirty three climate/economy "models" may potentially describe the impact of climate change on the economy. These models come from the combination of the eleven CCR values describing the possible relationship between GHG emissions and temperatures, and the three different relationships between temperature increases and economic damages (see Section 2.2). Each model con-

[24]The gross level of output and the abatement cost are also potentially uncertain. However, since we want to focus on the type of uncertainty described above, here we do not consider these additional sources of uncertainty.

tains an aleatory component. For each of these thirty three models m_θ that we index by θ, it is possible to compute the expected reward $R(E, \theta)$ associated with any emission policy.[25] We can then express these thirty three reward functions in monetary terms by computing the certainty equivalents $c(a, \theta) = u^{-1}(R(E, \theta))$. These certainty equivalents are represented in Figure 6.

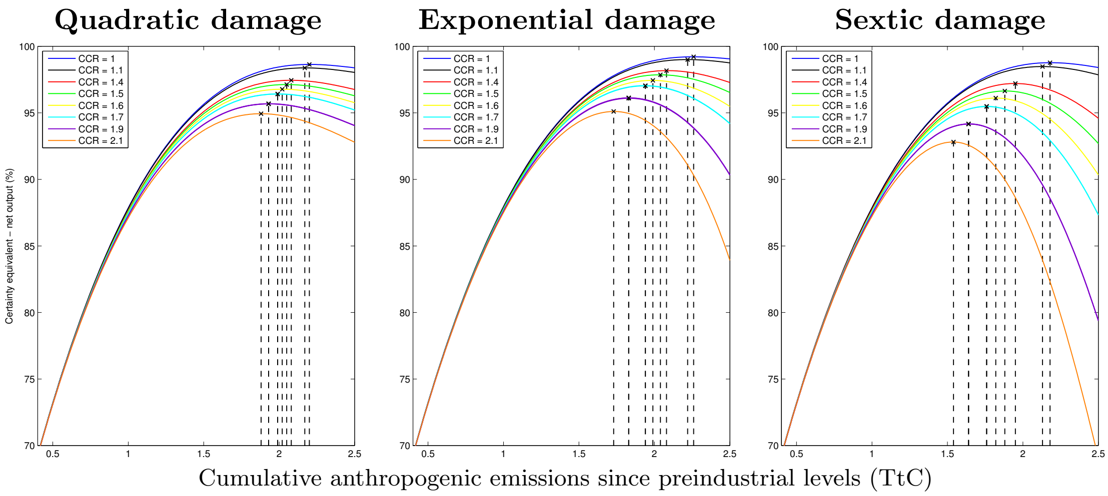

Figure 6: Certainty equivalents of output (net of climate damages and of abatement costs) as functions of cumulative emissions, for different CCR values and different damage functions: quadratic (1st column), exponential (2nd column) and sextic (3rd column)

The certainty equivalent represent, for each of the thirty three models and each level of cumulative anthropogenic emission, the certain amount of net output that a policy maker deems worth as much as the risky net output. In other words, each certainty equivalent represents a measure of net output that integrates the attitude towards the aleatory part of uncertainty. The policy maker, however, does not know which is the correct certainty equivalent. The certainty equivalent is, in that sense, itself uncertain since it depends on the values of the different structural parameters used. There are thus thirty three potential certainty equivalents, depending on whether the damage function is quadratic (first column), exponential (second column) or sextic (third column), and on the value of the CCR parameter (represented by different colors in Figure 6). For each particular model representing the impact of climate policy on economic output, it is possible to determine the optimal action to put in place. This is achieved by finding the level of cumulative

[25] In this example, the consequence function is simply the net output computed as $\rho(E, s) = \frac{Y_{\text{gross}} - C(E)}{1 + D(E,s)}$, where Y_{gross} is the gross output, $D(E, s)$ represents the damages associated with climate change, and $C(E)$ is the abatement cost. Both damages and costs depend on the action taken (E represents the level of GHG emissions). The damage function is uncertain as represented by the presence of s, encompassing both a random and a structural components. The abatement cost function is assumed to be nearly cubic as in Nordhaus and Sztorc (2013). The von Neumann-Morgernstern utility function u used is a power function, with a constant relative risk aversion coefficient of 1.5.

emissions maximizing the certainty equivalents (represented by the dashed vertical lines). These optimal levels of cumulative GHG emissions since preindustrial levels range from 1.54 TtC to 2.26 TtC, depending on the model considered. Unsurprisingly, lower carbon-climate response parameters induce higher optimal levels of emissions, while the use of a sextic damage function to characterize the impact of climate change tends to favor lower emission policies.

Since the ranking of the certainty equivalents is the same as that of the expected utilities,[26] we can first analyze the results of Figure 6 in the light of the unanimity criterion. As can be observed, up to the level of 1.54 TtC any level of cumulative emissions is unanimously better than any inferior level of emissions. To see this, consider for example the climate policy aiming at 1.5 TtC of cumulative emissions. For any of the thirty three models presented in Figure 6, it might be checked that this policy dominates –i.e., it leads to a higher level of certainty equivalent– any other policy with a lower level of cumulative emissions (say for example 1 TtC). Analogously, for emission levels superior to 2.26 TtC, a policy aiming at a specific level of cumulative emissions is always unanimously dominated by any policy aiming at a lower level. For levels of cumulative emissions between these thresholds 1.54-2.26 TtC, it is however impossible to find any policy satisfying the unanimity condition. The incompleteness of unanimity preferences therefore prevents any decision to be taken in such a situation, so other decision criteria need to be followed if the policy maker has to make a choice.

If the policy maker decides to behave extremely precautionary by taking into account only the model giving her the lowest expected reward, she only considers the combination CCR= 2.1/sextic damage, and fixes the level of cumulative emissions to 1.54 TtC. This policy maker is extremely uncertainty averse in that she uses Wald's maxmin criterion (13) illustrated in black in Figure 7. Alternatively, if the policy maker considers aleatory and model uncertainty the same way, she aggregates the expected rewards by taking a weighted average over them, where the weights represent her degree of belief in each specific model. In practice, this means that an overall certainty equivalent aggregating the different certainty equivalents associated with each specific model may be computed. This overall certainty equivalent incorporates the policy maker's attitude towards model uncertainty exactly in the same way as it incorporates her attitude towards risk. The overall certainty equivalent under a uniform prior over the possible models $-\mu(\theta) = 1/33$ for all θ– is represented in blue in Figure 7. The decision criterion in this case is the classical subjective expected utility (SEU) criterion (9). The optimal decision is a cumulative level of emissions of 1.85 TtC since preindustrial levels, which corresponds to the solution of problem (10).

Instead, if the policy maker is averse to uncertainty, so dislikes more epistemic uncer-

[26] A certainty equivalent is nothing but a monotonic transformation of an expected utility.

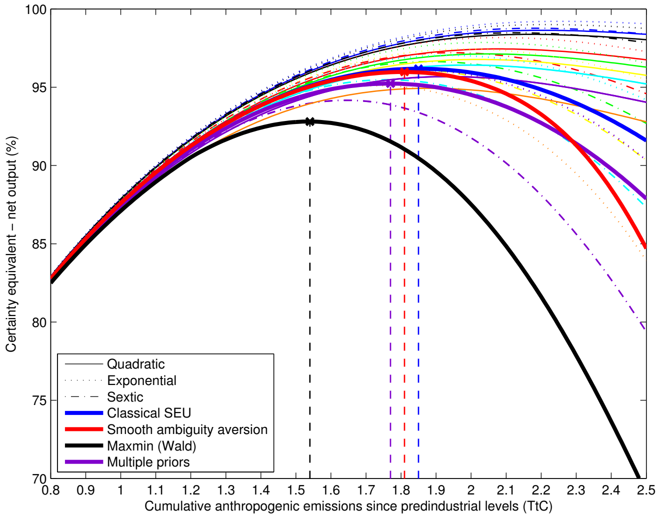

Figure 7: Decision making criteria and optimal decisions under uncertainty

tainty than risk, but is not as precautionary as a maxmin policy maker, she may compute an overall certainty equivalent by means of a function v –more concave than function u–representing her attitude towards model uncertainty. An example of such overall cer-

Criterion used	Optimal decision [in cumulative emissions since predindustrial levels (TtC)]
V_{mxm}	$a^*_{mxm} = 1.54$
V_{eu}	$a^*_{eu} = 1.85$
V_{smt}	$a^*_{smt} = 1.81$
V_{mp}	$a^*_{mp} = 1.77$

Table 3: Example of optimal policies under uncertainty

tainty equivalent is represented in red in Figure 7.[27] In this case, the decision criterion is the smooth one (14), and the optimal level of cumulative emissions is lower than under expected utility. It approximately corresponds to 1.81 TtC since preindustrial levels.

Finally, if the policy maker follows the multiple priors approach (15), she considers different probability measures over the models, compute the expected utility for each of them, and considers only the one giving her the lowest level of expected reward. An example of such overall certainty equivalent is represented in purple in Figure 7. It represents the minimum expected reward obtained for two distinct priors: the uniform one where all the thirty three models are weighted equally, and one which considers the lower values of CCR as implausible (and therefore puts a weight 0 to them and a uniform prior over the

[27]In this example the model uncertainty aversion function v is also a power function, with a constant relative model uncertainty aversion coefficient of 15. The prior distribution remains the uniform one.

remaining models). The optimal level of cumulative emissions under the multiple priors model in this situation is lower than under the expected utility one. It corresponds to 1.77 TtC since preindustrial levels. The optimal decisions for each of these criteria is summarized in Table 3.

8 Discussions and conclusion

While the application we presented in the previous section may be too simplistic to be actually used by policy makers to design climate policies, it enables us to illustrate the differences between alternative decision criteria that can be used in the presence of uncertainty.

When it comes to making a choice in the presence of uncertainty, which is then the criterion one should adopt? As argued above, the standard line of reasoning that has traditionally been followed in climate policy is that, to make coherent choices, the policy maker should use either von Neumann-Morgenstern expected utility if she knows the true model or, otherwise, Savage's subjective expected utility with respect to her subjective probabilities over alternative models. In the context of our example, what this approach implicitly assumes is that, if epistemic uncertainty is present, it is treated in the same way as aleatory uncertainty, so the policy maker uses the "predictive" model (11). For a long time, the expected utility theory has been seen as the only convincing approach to make rational choices under uncertainty. Following this idea, Broome (2012) for example writes that, in the context of policy decisions regarding global warming:

> "The lack of firm probabilities is not a reason to give up expected value theory. You might despair and adopt some other way of coping with uncertainty; you might adopt some version of the precautionary principle, say. That would be a mistake. Stick with expected value theory, since it is very well founded, and do your best with probabilities and values."

While the axiomatic foundations of the expected utility approach appear at a first glance compelling, the claim that they constitute a necessary condition for rationality in decision making has, however, been challenged at least since Ellsberg (1961). For example, recently Gilboa et al. (2008, 2009, 2012); Gilboa and Marinacci (2013) argue that behaving in accordance with Savage's axioms raises several difficulties and that relaxing the assumption that decision makers are Bayesian might well be rational. It does not mean that decision makers are unable to think probabilistically or fail to compute probabilities correctly, but rather that they acknowledge that expected utility requires more information than they actually have, so its use would require some arbitrary assumptions that supplement the limited information.

The decision frameworks that we present in this paper are consistent with such an interpretation of rationality, so they are compatible with a normative assessment of optimal climate policies. In a context in which a variety of alternative models exist –each implying a different stochastic forecast– but where information about the accuracy of each of them is limited, such alternative decision frameworks may prove to be desirable. Indeed, the decision adopted is *robust* in the sense that the selected action does reasonably well across a range of models (Mukerji, 2009). This property seems particularly valuable when the consequences of the actions taken have long-lasting and global impacts, as it is the case with climate change.

When being asked to take actions under uncertainty, policy makers might well feel it is more rational to use these alternative criteria than to follow a standard expected utility approach. All too often, uncertainty has been used as an excuse for insufficient action in climate policy making. The most important –but potentially most difficult– thing to do is to acknowledge that there are things which we just do not know and then act in consequence.

References

Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica 19*, 404–437.

Athanassoglou, S. and A. Xepapadeas (2012). Pollution control with uncertain stock dynamics: when, and how, to be precautious. *Journal of Environmental Economics and Management 63*(3), 304–320.

Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and T. Sargent (2016). Monetary policies in self-confirming equilibria with uncertain models. *IGIER Working Papers series, Bocconi University* (WP573).

Battigalli, P., A. Francethich, G. Lanzani, and M. Marinacci (2017). Learning and self-confirming long-run biases. *IGIER Working Papers series, Bocconi University* (WP588).

Beck, M. and T. Krueger (2016). The epistemic, ethical, and political dimensions of uncertainty in integrated assessment modeling. *Wiley Interdisciplinary Reviews: Climate Change 7*(5), 627–645.

Berger, L. (2016). The impact of ambiguity and prudence on prevention decisions. *Theory and Decision 80*(3), 389–409.

Berger, L. and V. Bosetti (2017). Are policymakers ambiguity averse? *mimeo*.

Berger, L., J. Emmerling, and M. Tavoni (2017). Managing catastrophic climate risks under model uncertainty aversion. *Management Science 63*(3), 749–765.

Bewley, T. F. (2002). Knightian decision theory. part I. *Decisions in economics and finance 25*(2), 79–110.

Bosetti, V., E. Weber, L. Berger, D. V. Budescu, N. Liu, and M. Tavoni (2017). COP21 climate negotiators' responses to climate model forecasts. *Nature Climate Change 7*(3).

Bradley, R., C. Helgeson, and B. Hill (2017). Climate change assessments: Confidence, probability, and decision. *Philosophy of Science 84*(3), 500–522.

Brock, W. A. and L. P. Hansen (2017). Wrestling with uncertainty in climate economic models. mimeo.

Broome, J. (2012). *Climate Matters: Ethics in a Warming World (Norton Global Ethics Series)*. WW Norton & Company.

Burke, M., M. Craxton, C. Kolstad, C. Onda, H. Allcott, E. Baker, L. Barrage, R. Carson, K. Gillingham, J. Graff-Zivin, et al. (2016). Opportunities for advances in climate change economics. *Science 352*(6283), 292–293.

Cerreia-Vioglio, S. (2016). Objective rationality and uncertainty averse preferences. *Theoretical Economics 11*(2), 523–545.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011). Uncertainty averse preferences. *Journal of Economic Theory 146*(4), 1275–1330.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013). Classical subjective expected utility. *Proceedings of the National Academy of Sciences 110*(17), 6754–6759.

Chambers, R. G. and T. Melkonyan (2017). Ambiguity, reasoned determination, and climate-change policy. *Journal of Environmental Economics and Management 81*, 74–92.

Convery, F. J. and G. Wagner (2015). Managing uncertain climates: Some guidance for policy makers and researchers. *Review of Environmental Economics and Policy*, rev003.

Deser, C., A. Phillips, V. Bourdette, and H. Teng (2012). Uncertainty in climate change projections: the role of internal variability. *Climate Dynamics 38*(3), 527–546.

Drouet, L., V. Bosetti, and M. Tavoni (2015). Selection of climate policies under the uncertainties in the fifth assessment report of the IPCC. *Nature Climate Change 5*, 937–940.

Dupuis, P. and R. S. Ellis (1997). *A weak convergence approach to the theory of large deviations*, Volume 902. John Wiley & Sons, New York.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics 75*, 643–669.

Friedlingstein, P., P. Cox, R. Betts, L. Bopp, W. Von Bloh, V. Brovkin, P. Cadule, S. Doney, M. Eby, I. Fung, et al. (2006). Climate-carbon cycle feedback analysis: Results from the C4MIP model intercomparison. *Journal of Climate 19*(14), 3337–3353.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory 118*(2), 133–173.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler (2010). Objective and subjective rationality in a multiple prior model. *Econometrica 78*(2), 755–770.

Gilboa, I. and M. Marinacci (2013). Ambiguity and the bayesian paradigm. In *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society.* D. Acemoglu, M. Arellano, and E. Dekel (Eds.). New York: Cambridge University Press.

Gilboa, I., A. Postlewaite, and D. Schmeidler (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy 25*(03), 285–296.

Gilboa, I., A. Postlewaite, and D. Schmeidler (2012). Rationality of belief or: why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese 187*(1), 11–31.

Gilboa, I., A. W. Postlewaite, and D. Schmeidler (2008). Probability and uncertainty in economic modeling. *The Journal of Economic Perspectives 22*(3), 173–188.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics 18*(2), 141–154.

Hansen, J., M. Sato, R. Ruedy, K. Lo, D. W. Lea, and M. Medina-Elizade (2006). Global temperature change. *Proceedings of the National Academy of Sciences 103*(39), 14288–14293.

Hansen, L. P. and M. Marinacci (2016). Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science 31*, 511–515.

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *The American Economic Review 91*(2), 60–66.

Hansen, L. P. and T. J. Sargent (2007). Recursive robust estimation and control without commitment. *Journal of Economic Theory 136*(1), 1–27.

Hansen, L. P. and T. J. Sargent (2008). *Robustness.* Princeton University Press.

Heal, G. and A. Millner (2014). Uncertainty and decision in climate change economics. *Review of Environmental Economics and Policy 8*(1), 120–137.

Hill, B. (2013). Confidence and decision. *Games and economic behavior 82*, 675–692.

Howard, P. (2014). Omitted damages: What's missing from the social cost of carbon. The cost of carbon project.

Howard, P., T. Sterner, et al. (2014). Loaded dice: Refining the meta-analysis approach to calibrating climate damage functions. In *2014 Annual Meeting, Minneapolis, Minnesota*, Number 169952. Agricultural and Applied Economics Association.

IPCC (2007). *Climate Change 2007: Mitigation.* [B. Metz, O. R. Davidson, P.R. Bosch, R. Dave, L. A. Meyer (eds.)] Cambridge, UK, and New York: Cambridge University Press.

IPCC (2013). *Climate Change 2013: The Physical Science Basis.* [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]. Cambridge, UK and New York: Cambridge University Press, 1535 pp.

IPCC (2014a). *Climate Change 2014: Impacts, Adaptation, and Vulnerability.* [Field, C.B., V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T.E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, L.L. White]. New York: Cambridge University Press, 1132 pp.

IPCC (2014b). *Climate Change 2014: Synthesis Report.* Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)]. IPCC, Geneva, Switzerland, 151 pp.

Keynes, J. M. (1921). *A Treatise on Probability.* Macmillan and Co, London.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*, 1849–1892.

Knight, F. (1921). *Risk, Uncertainty, and Profit.* Houghton Mifflin.

Koopmans, T. C. (1947). Measurement without theory. *The Review of Economics and Statistics 29*(3), 161–172.

Kunreuther, H., G. Heal, M. Allen, O. Edenhofer, C. B. Field, and G. Yohe (2013). Risk management and climate change. *Nature Climate Change 3*(5), 447–450.

Lange, A. and 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.

Lemoine, D. and C. P. Traeger (2016). Ambiguous tipping points. *Journal of Economic Behavior & Organization 132*, 5–18.

Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management science 52*(4), 514–528.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Mäki, U. (2011). Models and the locus of their truth. *Synthese 180*(1), 47–63.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association 13*(6), 1022–1100.

Marschak, J. (1953). Economic measurements for policy and prediction. In *Studies in Econometric Method* (W. Hood and T. J. Koopmans, eds.), pp. 1–26. Wiley, New York.

Marschak, J. and R. Radner (1972). *Economic theory of teams.* Yale University Press, New Haven.

Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009). The proportionality of global warming to cumulative carbon emissions. *Nature 459*(7248), 829–832.

Meinshausen, M., N. Meinshausen, W. Hare, S. C. Raper, K. Frieler, R. Knutti, D. J. Frame, and M. R. Allen (2009). Greenhouse-gas emission targets for limiting global warming to 2 c. *Nature 458*(7242), 1158–1162.

Millner, A., S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics 55*(1), 21–46.

Mukerji, S. (2009). Foundations of ambiguity and economic modelling. *Economics and Philosophy 25*(03), 297–302.

Nordhaus, W. and P. Sztorc (2013). Dice 2013r: Introduction and user's manual. *retrieved November*.

Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the" dice" model. *The American Economic Review*, 313–317.

Pindyck, R. (2007). Uncertainty in environmental economics. *Review of Environmental Economics and Policy 1*(1), 45.

Pindyck, R. S. (2013a). Climate change policy: What do the models tell us? *Journal of Economic Literature 51*(3), 860–72.

Pindyck, R. S. (2013b). The climate policy dilemma. *Review of Environmental Economics and Policy 7*(2), 219–237.

Pindyck, R. S. (2015). The use and misuse of models for climate policy. Technical report, National Bureau of Economic Research.

Rezai, A. and R. van der Ploeg (2017). Climate policies under climate model uncertainty: max-min and min-max regret. *Energy Economics* (forthcoming).

Rudik, I. (2016). Optimal climate policy when damages are unknown.

Savage, L. (1954). *The Foundations of Statistics.* New York: J. Wiley. second revised edition, 1972.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*(3), 571–587.

Stern, N. et al. (2016). Current climate models are grossly misleading. *Nature 530*, 407–409.

Van den Bergh, J. C. (2004). Optimal climate policy is a utopia: from quantitative to qualitative cost-benefit analysis. *Ecological economics 48*(4), 385–393.

von Neumann, J. and O. Morgenstern (1947). *Theory of Games and Economic Behavior.* 2nd ed. Princeton University Press.

Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity.* Cambridge university press.

Wald, A. (1950). *Statistical decision functions.* New York: John Wiley & Sons.

Weitzman, M. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics 91*(1), 1–19.

Weitzman, M. L. (2012). GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory 14*(2), 221–244. 00066.

Xepapadeas, A. and A. Yannacopoulos (2017). Spatially structured deep uncertainty, robust control, and climate change policies. mimeo.