DEPARTMENT OF INTERNATIONAL AND EUROPEAN ECONOMIC STUDIES

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Alternative Types of Ambiguity and their Effects on the Probabilistic Properties and Tail Risks of Environmental-Policy Variables

Phoebe Koundouri

Nikitas Pittis

Panagiotis Samartzis

Nikolaos Englezos

Andreas Papandreou

Working Paper Series

17-03

August 2017

Alternative Types of Ambiguity and their Effects on the Probabilistic Properties and Tail Risks of Environmental-Policy Variables

Phoebe Koundouri[*] Nikitas Pittis[†] Panagiotis Samartzis[‡]

Nikolaos Englezos[§] Andreas Papandreou[¶]

Abstract

The concept of ambiguity with respect to decision making about climate change has recently attracted a lot of research interest. The standard approach for introducing ambiguity into this framework is to assume that the decision maker (DM) exhibits ambiguity aversion, with the latter being represented by axioms on DMs preferences different than Savage's (*sure-thing principle*). As a result, DM is deprived of the property of probabilistic sophistication, since she is faced with either multiple prior probability functions, or a single but incoherent one (*capacity*). This paper approaches the issue of ambiguity with respect to climate change from a different perspective. In particular, we assume that ambiguity does exists but it does not affect the formation of DMs prior probability function. Instead, it affects the formation of her posterior probability function. Specifically, we assume that there are n experts, who supply DM with probabilistic input. Hence, although DM has a well defined prior (formed before any expert information on objective probabilities has arrived), she cannot decide which piece of information should conditionalize upon (*defer to*). We refer to this type of ambiguity as "deferential ambiguity" and show that it affects both DM and the experts. We also introduce a second type of ambiguity, which is solely born by the experts. This type of ambiguity stems from the experts potential inability to discern DMs preferences. This ambiguity is referred to as "preferential ambiguity" in the paper. The main objective of the paper is to analyze the possible interactions between the two types of ambiguity mentioned above and to assess their impact on the probabilistic properties (in particular, tail risks) of environmental-policy variables.

[*]School of Economics, Athens University of Economics and Business, Greece; Grantham Research Institute, London Schoopl of Economics and Political Science, UK; International Center for Research on the Environment and the Economy, Greece; e-mail: pkoundouri@aueb.gr (corresponding author)

[†]Department of Banking and Financial Management, University of Piraeus, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: npittis@unipi.gr

[‡]Department of Banking and Financial Management, University of Piraeus, Greece; School of Economics, Athens University of Economics and Business, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: psamartzis@unipi.gr

[§]Department of Banking and Financial Management, University of Piraeus, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: englezos@unipi.gr

[¶]School of Economics, National and Kapodistrian University of Athens, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: aap@econ.uoa.gr

Keywords: decision making on climate change, ambiguity, deep uncertainty, Savage's sure-thing principle, deferential ambiguity, preferential ambiguity, tail risks of environmental-policy variables.

JEL Classification: D8, D80, D81, D83, D84, D89

1 Introduction

It is often claimed that decision making on climate change is characterized by *ambiguity* (or deep uncertainty). The term "ambiguity" is used, in a practical sense, to portray the massive uncertainty that besets the effects of emissions of greenhouse gases on global warming and in a more formal sense, to characterize the epistemic state of the agents involved in the determination and evolution of the climate-related variable, say Y. To this end, we may distinguish three types of such agents. (a) The decision maker (DM), whose actions today affect the level of Y tomorrow. (ii) The expert (or experts) who either knows the true model for the determination of Y over time or he possesses an approximation of it and (iii) the public. Ambiguity, being an epistemic state rather than a physical property may affect some or all of the agents mentioned above. The main aim of the paper is first, to characterize the ambiguity faced by DM as well as the one faced by the expert, second to examine their interactions and third to analyze their effects on the probabilistic properties of Y.

More specifically, the problem of ambiguity that arises in the context of decision making akin to climate change may be described as follows: Consider a decision maker, DM, who at time t is about to form her system of probabilistic beliefs, that is her subjective probability function, P_t^{DM}, defined on a field of propositions/events Σ. DM is assumed to be rational, which amounts to saying that (i) DM's subjective probability function is coherent for every t, (ii) DM updates her probabilistic beliefs in the light of new evidence by Bayesian conditionalization (BC) and, (iii) DM obeys the Principal Principle (PP) (see, Lewis 1980). PP (roughly) states that if DM knows the objective probability (chance), $Ch_t(A)$, of $A \in \Sigma$, then she sets her subjective probability of A equal to the corresponding objective probability.

Let us now bring another type of agent into the picture, namely the expert(s). An expert in this context is defined to be the agent who knows the objective probabilities (chances), $Ch_t(A)$, $A \in \Sigma$ of the events of interest. In the case that there is a unique expert, DM is most likely to perceive this expert as the true bearer of objective probabilities, which in turn implies that she will have a strong incentive (especially if DM is benevolent) to defer to him at each point in time. As a result, DM's subjective probability distribution always coincides with the corresponding (unique) objective probability distribution. Hence, DM always knows the true probabilities of the events of interest, which in turn implies that she always operates under an environment of "risk" (known probabilities) rather than "ambiguity" (unknown probabilities).

Ambiguity arises in the case that there are more than one experts, say n, who disagree with each other about the chances of the events in Σ. Equivalently, one may assume that there are n competing statistical models for the same phenomenon. Each of these experts has his own view about the "objective" probability function on Σ (his own model). Hence, DM is faced with n "objective" probability functions, $P_{i,t}^{X}$, instead of one, which in turn complicates her attempts to form her subjective probability function P_t^{DM}. How can one interpret and model these complications?

To this end, there are two options. What distinguishes the one option from the other is the time at which DM is assumed to form her prior probability function (or, her initial credence function). The first option, hereafter referred to as Classical Bayesianism (CB) (see Meecham 2007), assumes that DM forms her prior probability function at t. This assumption implies that DM's current probability function, P_t^{DM} is entitled to be treated as the "prior". In such a case, DM is allowed to use all the information that is available to him at t, that is both background and specific information, in setting her prior. The second option, hereafter referred to as Modern Bayesianism (MB) (see Meecham 2007), insists that only background information is allowed to affect the formation of DM's prior probability distribution. Specific information, such as information on objective chances supplied by the experts, should be treated strictly as "conditioning information" that affects DM's probabilistic beliefs at t strictly through BC. Easwaran and Fitelson (2011) define such a prior probability function (initial credence function) as one "...which would not be informed by specific bodies of empirical knowledge regarding objective chances." (2011, pp. 4). This in turn implies that DM's current probability function, P_t^{DM}, cannot be treated as prior; instead the prior, P_0^{DM}, was determined at a time "prior" to t, say $t = 0$, in which only background information was available.

It is worth noting that despite terminology that suggests otherwise, MB goes at least as far back as Carnap (1950). Indeed, his monumental work on logical probability and induction is based on the concept of hypothetical or counterfactual initial credence function that can be ascribed to an agent, before the collection of any evidence (including information on objective probabilities). According to Carnap, P_0^{DM} reflects the agent's permanent dispositions for forming beliefs, as opposed to P_t^{DM} which reveals merely her momentary inclinations at time t.

The aforementioned distinction between CB and MB has serious implications for the way in which ambiguity is defined, accounted for and finally modelled. The so-called "ambiguity aversion" literature adopts the CB interpretation in the context of which DM, in the face of contradicting (or ambiguous) information, cannot decide how to form a proper P_t^{DM}. As a result she may end up adopting a set of priors (Gilboa and Schmeidler 1989), or alternatively, a single non-additive prior (capacity) (Schmeidler 1989), thus being deprived of the virtue of probabilistic sophistication and failing to satisfy Savage's axioms. To this end, alternative, more permissive axiomatizations have been suggested in an attempt to rationalize DM's predicament (for a critical survey of this literature, see Al-Najjar and Weinstein 2009). A common feature of these alternative axiomatic systems is that they all relax Savage's *sure thing principle*, while at the same time introduce axioms that represent some form of ambiguity-sensitive behaviour. The interesting aspect of these systems from an empirical point of view is that they yield new criteria for decision making under uncertainty (or, rather, under ambiguity), which are different from the classical "maximization of subjective expected utility". Heel and Millner (2015) examine how the adoption of such criteria (specifically, the MaxMin and Smooth Ambiguity ones put forward by Gilboa and Schmeidler 1989 and Klibanoff, Marinacci and Mukherji 2005, respectively) affect decision making in the context of climate policy.

Adopting the MB perspective offers a radically different interpretation of DM's ambiguity. Here, the multiplicity of "objective" distributions (supplied by the experts at t) does not threat the formation of DM's prior. This is because, DM's prior has already been formed at a period prior to t (in our case at t=0) in which no specific information was available. Put differently, at the time that the n experts furnish their views, DM is already equipped with her prior P_0^{DM}, so that this prior is immune to the ambiguity of information regarding objective probabilities. These objective probabilities (or any other type of probabilistic information such as point forecasts) are treated as "data" or "information", upon which DM conditionalizes (using P_0^{DM} as vehicle) in order to update her beliefs from P_0^{DM} to P_t^{DM}. Despite its theoretical elegance this approach, usually referred to as *Supra Bayesian* method (see, for example, French 1985, Lindley 1985, Jacobs 1995), is not easy to apply in some real-world situations (see Jacobs 1995 for a discussion of these difficulties).

Another method for combining or aggregating experts's probabilistic input is the so called *axiomatic method* which does not make an explicit use of DM's prior probability function. Instead, this method is based on (i) setting a number of desirable axioms that the combined distribution should satisfy and (ii) find the functional form that satisfies most (if not all) of these axioms. One of the most widely used functional form is the so called linear opinion pool, according to which

$$P_t^{DM} = \sum_{i=1}^{n} w_i P_{i,t}^{X}, \tag{1}$$

where the weights w_i are non-negative and sum to one. As Clemen and Winkler (1999) remark, the weights w_i may be interpreted as representing the relative quality of the n experts. In the case that all the experts are regarded as equivalent (by DM), (1) reduces to a simple arithmetic average. The linear opinion pool satisfies the axioms of *unanimity* and *marginalization* (see, for example Clemen and Winkler 1999). However, it fails to satisfy the principle of *External Bayesianity*, which is one of the reasons why an alternative combination scheme, the so-called logarithmic opinion pool, is occasionally employed.

In all the cases presented above, there seems to be a common underlying assumption concerning the relationship between (i) the probabilistic input about the phenomenon Y that DM receives from the experts (ii) her actions based on this input and (iii) the actual probabilistic properties of Y. In particular, it is tacitly assumed that DM's actions (informed by the views of the experts) do not affect the actual probability distribution of Y, or put differently, DM's actions are exogenous to Y. Instead, it is assumed that DM requires at $t-1$ the experts's probabilistic forecasts for Y_t, in order to form her own probabilistic assessment at $t-1$ for Y_t. Then, under the resulting set of probabilistic beliefs, DM will take the (perceived as) optimal course of action at t-1 which, however, does not affect the generation of Y_t at t. As a result, the experts do not have to anticipate the DM's actions before they elicit their probabilistic views on Y_t because they do not have to account for these actions in their models. In such a case, we have only one-way causality between the probabilistic views of DM and those of the experts: Experts' views affect DM's views (via PP), but not the other way round.

This assumption, however, does not seem to be a realistic one concerning issues of climate change. For example, Heel and Millner (2015) comment on the endogeneity of emissions as follows: "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." (pp. 5). This is a case in which forecasts result in an adaptive change which in turn affects the forecasted quantity. Consequently, an expert who tries to produce a forecast for the change, ΔY_t, in emissions between $t-1$ and t is forced to consider, as one of the causal factor in his model, the DM's forecast for next period's emissions. Why? Because these forecasts will determine DM's course of action, which in turn is likely to affect the actual change of emissions. This feature produces a two-way causality between the probabilistic views of DM and those of the experts: Experts' forecasts of ΔY_t affect DM's forecasts of ΔY_t, but at the same time, DM's forecasts, being causal factors in experts's models for ΔY_t, affect experts' forecasts.

The main aim of the paper is to investigate the effects from the interaction between DM's and experts' forecasts for ΔY_t on the actual (objective) distribution, D_Y, of ΔY_t. Various forms of such interactions are analyzed, with each one generating a different level of ambiguity in either DM and/or the experts. The source of DM's ambiguity is the potential disagreement among the experts (from now one we assume that there are at most two experts). If such disagreement exists, then DM is not certain to which expert should defer, or more generally, how she should combine their views.

Concerning experts' ambiguity, we investigate the following two sources: (i) One source of experts' ambiguity is the aforementioned uncertain deferential attitude of DM. Specifically, at each point in time, expert $i, i = 1, 2$ faces the following possibilities: (a) DM defers to his own forecasts (forecasts of i). (b) DM defers to the forecasts of expert $j \neq i$ (c) DM defers to a combination (e.g. linear pooling) of the two experts' forecasts (d) DM defers to none of the two. These possibilities raise for each of the experts the following "specification issue": how should DM's deferential attitude be introduced in each of the experts' models? The way that each expert answers this question bears different implication for the actual generation mechanism of ΔY_t. Hereafter, we shall refer to this type of ambiguity as "deferential ambiguity". (ii) A second source of experts' ambiguity (present even in the case of a single expert) stems from the potential inability of DM to correctly discern the society's preferences about the desired change of Y_t at each point in time. Hereafter, this type of ambiguity will be referred to as "preferential ambiguity". As will be shown below, even a DM with "simulated attitude" (Tunney and Ziegler 2015) that, is one who always tries to act according to society's preferences, she has to identify correctly these preferences, in order for her actions to be successful. Assume that the true preferences of society formed at $t-1$ for the value of Y at t are denoted by $Y^*_{(t-1),t}$. More specifically, assume that the society determines $Y^*_{(t-1),t}$ as the sum of the current value, Y_{t-1}, of Y (limited degrees of freedom due to physical constraints) plus a quantity Z_{t-1} which represents the desired change in Y between $t-1$ and t, that is $Y^*_{(t-1),t} = Y_{t-1} + Z_{t-1}$. This rule of dynamic determination of social preferences may be justified by assuming that in forming its desired level of Y for next period, the society takes into account the current level of Y. Put differently, tomorrow's level of desired Y is (physically) constrained by the current level of Y. Assume that the expert is able to identify and measure correctly Z_{t-1}, which means that the expert knows accurately society's preferences at $t-1$. Let us further assume, that despite her best efforts, DM fails to diagnose Z_{t-1}, but instead she believes that society's preferences are best captured by W_{t-1}. Alternatively, she might be able to identify Z_{t-1}, but she believes that society is currently wrong in focusing on Z_{t-1}. In Tunney and Ziegler's (2015) taxonomy, such a DM exhibits "benevolent attitude", in the sense that her actions are driven not by identifying "what the society would do" but rather by "what the society should do". This possibility is akin to the so-called "centralism" thesis (Press 1994) according to which "ecological problems can be solved only by strong centralized control of human behaviour, thus making common resource decisions by central authorities and replacing democratic rule by 'ecological mandarins' with the 'esoteric' knowledge and public spirit required" (Coenen et. al 1998, pp5)).

Irrespective of the reasons that make DM to adopt W_{t-1} instead of Z_{t-1}, the crucial question is the following: Does the expert know that DM does not act upon Z_{t-1} but upon W_{t-1}? To this end there are three possibilities: (1) The expert (being a true expert) knows DM's preferential error right from the start. In such a case, the effects of DM's error on the actual generation process of ΔY_t are relatively simple to analyze: The probabilistic properties of ΔY_t do not depend on the probabilistic properties of the stochastic process $\{Z_t\}$, but rather on those of $\{W_t\}$. (2) The expert never realizes DM's error and erroneously believes that DM acts on Z_{t-1}. In this case, the actual distribution of ΔY_t will be different than the one in experts mind, which in turn implies that the expert is never the bearer of objective chance. (3) The expert initially believes that DM acts on Z_{t-1}, but he endorses a learning process, in the context of which he repeatedly compares the realized values of ΔY_t with those implied by his model. In general, the resulting asymptotic distribution of ΔY_t does not coincide to that in the expert's mind, since there is a non-zero bias that survives even asymptotically. As expected, the probabilistic properties of ΔY_t in this case differ significantly from those of the previous case.

Next section, defines our basic model and examines the benchmark case which is defined by the following assumptions: (i) DM is interested in experts' point forecasts (conditional expectations) rather than their views about the conditional or unconditional distribution of ΔY_t. (ii) DM is assumed to be a "projectivist" one, who always (i.e. for each t) complies with society's preferences by acting in such a way as to bring the actual Y_t in line with the level $Y^*_{(t-1),t}$ designating by society at t-1 as optimal for t. (iii) DM acts upon Z_{t-1} rather than W_{t-1}, which means that DM's perception about the optimal change in Y_t coincides with that of the society. In such a case, her projectivist attitude can be fulfilled. (iv) there is only one expert who knows the structural form of the statistical model describing the probabilistic properties of $\{\Delta Y_t\}$, as well as the true values of the model's structural parameters. As a by product of this assumption, the expert knows that DM acts on the basis of Z (rather than W). Obviously, in our benchmark case, there is neither deferential not preferential ambiguity. As such, this case is equivalent to the basic case of Baillon, Cabantous and Wakker (2012) (referred to as the "source risk") according to which "both agents are Bayesian and agree with each other (and everyone else). This is the common case of generally accepted objective probabilities, with no ambiguity involved" (pp.116). However, the aforementioned authors do not allow for any interactions between DM and the expert: "We also assume that there is no interaction between the agents themselves, or between the agents and the decision maker, so that no group process is involved." (pp. 116). On the contrary, our study not only allows for such interactions but makes them our central topic of research.

Section III introduces preferential ambiguity. This is carried out by retaining the first two assumptions of the foregoing benchmark case, and replacing (iii) and (iv) by (iiia) and (iva) respectively: (iiia) DM acts upon W_{t-1} rather than Z_{t-1}. (iva) The unique expert knows all the features the structural form of the model of $\{\Delta Y_t\}$, but one, namely the fact that DM employs W_{t-1} instead of Z_{t-1}. This means that the expert will produce his forecasts of ΔY_t by means of a misspecified model which includes the wrong variable Z_{t-1} instead of the true one, W_{t-1}. In this section we also assume that the expert never learns about his specification error, (no learning mechanism is assumed to be in place), which in turn implies that the expert ends up having a subjective probability of ΔY_t different than the objective one (the expert is not the bearer of objective chance).

Section IV relaxes the no-learning assumption of Section III, and derives the asymptotic distribution of ΔY_t under recursive OLS learning. An interesting feature of this case is that the forecast error committed by the expert never goes to zero. However, even under the aforementioned asymptotic bias, the stochastic process $\{\Delta Y_t\}$ converges-in-law. Section V gives a brief description of how deferential uncertainty can be introduced in the model, together with its potential interactions with preferential uncertainty and their combined effects on the probabilistic properties of $\{\Delta Y_t\}$. Second VI summarizes the main findings, draws lines for future research and concludes the paper.

2 Basic Model and Assumptions (No Deferential Uncertainty)

Let us begin with introducing some basic concepts and notation (with the risk of repeating some of those already introduced in the previous section). Assume that $\mathcal{E}_{t-1}(Y_t)$ denotes the DM's forecast (made at $t-1$) for the actual level of Y at t, whereas Y^*_{t-1} (a simplified notation for $Y^*_{(t-1),t}$) stands for the level of Y that the society at $t-1$ thinks of as optimal (or desired) at t. To make things more concrete, assume that Y represents the level of CO_2 emissions. As already mentioned, DM is supposed to act in line with society's preferences. This implies that whenever $\mathcal{E}_{t-1}(Y_t) > Y^*_{t-1}$, ($\mathcal{E}_{t-1}(Y_t) < Y^*_{t-1}$) DM acts in such a way as to produce a negative (positive) actual change ΔY_t. As far as Y^*_{t-1} is concerned, we assume that it is the sum of the actual Y_{t-1} and another variable Z_{t-1}, with the latter representing society's preferences at $t-1$ for the next period's value of Y. Specifically

$$Y^*_{t-1} = Y_{t-1} + Z_{t-1} \tag{2}$$

This means that at time $t-1$, the society sets the desired level for Y at time t, as the sum of the current level Y_{t-1} plus the adjustment factor Z_{t-1}. If $Z_{t-1} > 0$ ($Z_{t-1} < 0$), the society prefers (at t-1) a higher (lower) value of Y_t than the one that currently prevails (namely Y_{t-1}). Let

$\mu_Z = E(Z_t)$ and $\sigma_Z^2 = V(Z_t)$. The more the society's targets change over time, the larger the value of σ_Z^2. What are the reasons that cause the society to change the desired level of Y over time? One important such reason may be a change in the society's perception about the sensitivity of the climate to CO_2 emissions. Indeed, sensitivity estimates exhibit significant variability both within the same period and across time. Changes in society's perception about the sensitivity parameter induce changes in society's estimates about the economic costs of emissions. As McKitrick (2014) remarks: "In a low-sensitivity model, GHG (greenhouse gases) emissions lead only to minor changes in temperature, so the socioeconomic costs associated with the emissions are minimal. In a high-sensitivity model, large temperature changes would occur, so marginal economic damages of CO2 emissions are larger." (pp. 1). For example, in a period $t = t_1$ in which the economic cost of emissions is estimated to be small, the society is likely to set $Z_{t_1} > 0$.

Does DM adopt society's target Z_{t-1} or deviates from that? As already mentioned, this section assumes that the answer to this question is in the affirmative. Moreover, since DM is supposed to act in the best interests of the society, $(\mathcal{E}_{t-1}(Y_t) - Y_{t-1}^*)$ enters as a causal factor in the determination of ΔY_t with a negative coefficient. We may refer to this factor as the "human" factor. In addition, there is a physical variable (assumed to be exogenous in the standard sense) X_t, affecting ΔY_t, which may be referred to as the "physical" factor. Bringing these two factors together results in the following equation,

$$\Delta Y_t = Y_t - Y_{t-1} = \alpha(\mathcal{E}_{t-1}(Y_t) - Y_{t-1}^*) + \beta X_t, \tag{3}$$

The structural parameters α and β are assumed to be time invariant. With respect to α (key parameter in the ensuing analysis) we assume $-1 < \alpha < 0$ in order to capture DM's socially-sensitive behaviour. Concerning the exogenous variable, we assume for simplicity that X_t is a Gaussian IID process with zero mean,

$$X_t \sim NIID(0, \sigma_X^2). \tag{4}$$

Under the assumptions made thus far, (3) becomes

$$\Delta Y_t = \alpha(\mathcal{E}_{t-1}(\Delta Y_t) - Z_{t-1}) + \beta X_t \tag{5}$$

As far as the experts are concerned, we assume that there is only one expert who knows the structural model given by equations, (2), (3) and (4) together with the value of the parameter vector $\theta = [\alpha, \beta, \mu_Z, \sigma_Z^2, \sigma_X^2]$. This means that the expert (at $t-1$) specifies in his model the same variable Z_{t-1} that DM adopts (no preferential ambiguity). Furthermore, we assume that DM always defers to expert's point forecast, $E_{t-1}(\Delta Y_t)$, and the expert is aware of this fact (no deferential ambiguity).

Under the foregoing assumptions, the solution of the model is quite simple: Taking expectations on both sides of (5), we get

$$E_{t-1}(\Delta Y_t) = -\frac{\alpha}{1-\alpha} Z_{t-1}$$

and therefore:

$$\Delta Y_t = -\frac{\alpha}{1-\alpha} Z_{t-1} + \beta X_t. \tag{6}$$

The conditional distribution of ΔY_t is

$$\Delta Y_t \mid \mathcal{F}_{t-1} \sim N\left(-\frac{\alpha}{1-\alpha} Z_{t-1}, \beta^2 \sigma_X^2\right),$$

where $\mathcal{F}_{t-1}$ represents the information until $t-1$. The corresponding unconditional distribution is

$$\Delta Y_t \sim N\left(-\frac{\alpha}{1-\alpha} \mu_Z, \left(\frac{\alpha}{1-\alpha}\right)^2 \sigma_Z^2 + \beta^2 \sigma_X^2\right). \tag{7}$$

Comments

(i) Since the coefficient $-\frac{\alpha}{1-\alpha}$ is always positive, the unconditional mean of ΔY_t is positive (negative) whenever the mean, μ_Z, of Z is positive (negative). This means that if the society desires, on average, a positive (negative) change in next period's level of emissions, this desire will be translated (via DM's actions) into an actual average positive (change) change in emissions.

(ii) The results for the unconditional variance of ΔY_t bear different implications for the role of the human factor than those for the unconditional mean, discussed above. More specifically, the human involvement in the generation process of Y_t always results in an increase in the variability of ΔY_t, (compared to the case that the human intervention were completely absent, that is $\alpha = 0$ by a factor equal to $\left(\frac{\alpha}{1-\alpha}\right)^2 \sigma_Z^2$. Put differently, even if the society (almost) always desires a lower level of next period's emissions (that is $Z_{t-1} < 0$), the DM's actions to achieve this task will produce an increase in the volatility of the actual changes in emissions, compared to the cases that (i) DM is inactive ($\alpha = 0$) or (ii) DM is active ($\alpha < 0$) but the society does not change its preferences over time ($Z_{t-1} = 0$).

(iii) Ceteris paribus, as the algebraic value of α decreases (becomes "more negative" so to speak), or equivalently as DM becomes more active, the mean of ΔY_t increases for $\mu_Z > 0$ and decreases for $\mu_Z < 0$. On the other hand, the effect of such an decrease in α on the unconditional variance of ΔY_t is unambiguous: the lower the value of α, the higher the variance. This case bears some interesting implications regarding a potential DM's decision to become more active, that is to decrease the value of α from α_0 (say -0.2) to α_1 (say -0.8) on "tail risks" (that is, the probability of realization of a very large value of ΔY_t). Consider the case in which $\mu_Z < 0$, that is the case in which the society exhibits aversion to emissions. The rise in the degree of DM's response (from α_0 to α_1), on the one hand decreases the mean of ΔY_t but on the other hand increases the variance of ΔY_t, with the net effect of these two opposite forces on the right tail of the distribution of ΔY_t being ambiguous.

To put it differently, the probability of observing larger values, given on the one hand, that the mean of the distribution of ΔY_t is shifted to the left and on the other, that the variance increases, could either increase or decrease depending on the distribution of ΔY_t. Under the normality assumption, however, this ambiguity is eliminated: the probability of observing larger values decreases as the degree of DM's response increases (a decreases). The graph below, depicts the probability of the event $E = \{\Delta Y_t > 0\}$ as a function of α, assuming, without loss of generality, that $\mu_Z = -0.5, \sigma_Z^2 = \beta = \sigma_X^2 = 1$:

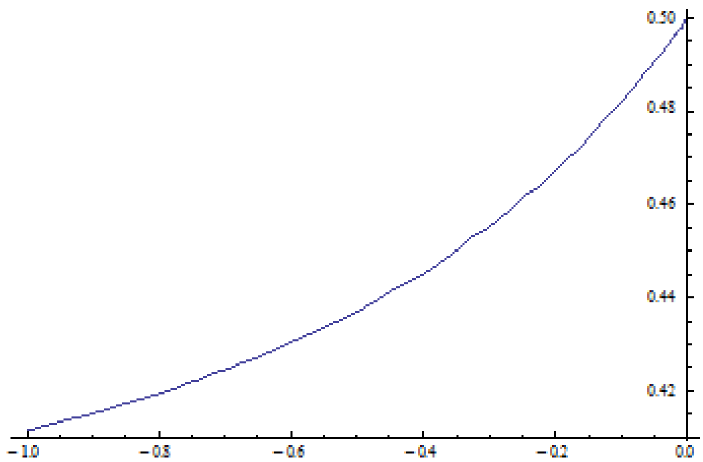

Figure 1: $P(E)$ as a function of α

3 Preferential Ambiguity with No Learning

Let us now relax the assumption that DM's and society's preferences are aligned. As mentioned above, DM acts upon her own preference variable W rather than that of the society. Her incentives of doing so may not be inferior. As Tunney and Ziegler (2015) remark, "surrogate decision makers may not have as their goal to match the wishes of the recipient, but instead to make what they perceive to be an optimal or benevolent decision." (2015, pp 884). Concerning W, let $\mu_W = E(W_t)$,

$\sigma_W^2 = V(W_t)$, $corr(W_t, Z_t) = \rho_{WZ}$ and $corr(W_t, X_{t+1}) = 0$. As a result, (3) becomes

$$\Delta Y_t = \alpha(\mathcal{E}_{t-1}(\Delta Y_t) - W_{t-1}) + \beta X_t \tag{8}$$

The expert fails to recognize the discrepancy between DM's and society's preferences. Hence, he believes that the law of motion of Y_t is given by

$$\Delta Y_t = \alpha(\mathcal{E}_{t-1}(\Delta Y_t) - Z_{t-1}) + \beta X_t. \tag{9}$$

He also believes (correctly) that $\mathcal{E}_{t-1}(Y_t) = E_{t-1}(Y_t)$ i.e. there is no deferential ambiguity. In this case, the expert's subjective conditional distribution (the one perceived by the expert as true) is given by

$$\Delta Y_t \mid \mathcal{F}_{t-1} \sim N\left(-\frac{\alpha}{1-\alpha} Z_{t-1}, \beta^2 \sigma_X^2\right) \tag{10}$$

Since DM always defers to the expert, it follows that

$$\Delta Y_t = -\frac{\alpha^2}{1-\alpha} Z_{t-1} - \alpha W_{t-1} + \beta X_t.$$

The objective conditional distribution of ΔY_t is

$$\Delta Y_t \mid \mathcal{F}_{t-1} \sim N\left(-a\left(\frac{\alpha}{1-\alpha} Z_{t-1} + W_{t-1}\right), \beta^2 \sigma_X^2\right), \tag{11}$$

whereas, the corresponding objective unconditional distribution is

$$\Delta Y_t \sim N\left(-\alpha\left(\frac{\alpha}{1-\alpha}\mu_Z + \mu_W\right), \left(\frac{\alpha^2}{1-\alpha}\right)^2 \sigma_Z^2 + \alpha^2 \sigma_W^2 + \beta^2 \sigma_X^2 + 2\frac{\alpha^3}{1-\alpha}\rho_{WZ}\sigma_W \sigma_Z\right). \tag{12}$$

Comments

(i) By comparing (10) with (11), we conclude that the variance of the expert's subjective distribution is equal to that of the objective conditional distribution, both being equal to $\beta^2 \sigma_X^2$. However, these two distributions are quite different with respect to their conditional means.

(ii) By comparing the unconditional mean in (12) with the corresponding one in (7), we observe that there are cases in which DM's focusing on W instead of Z, that is her "deviant behaviour", results in significant shifts in the unconditional distribution of Y_t, which in turn may prove beneficial for the society in the future. Specifically, if $\mu_W < -\frac{\alpha}{1-\alpha}\mu_Z$, then DM's actions shift the unconditional distribution to the left, thus reducing the probability of an extremely large value of ΔY_t (tail event) in the future. This may be interpreted as the result of DM's benevolent behaviour who acts on the basis of what the society should prefer at $t-1$ (normative stance) rather than what the society does prefer at $t-1$ (descriptive stance).

(iii) A similar observation can be made with respect to the unconditional variance. In particular, if $\sigma_Z > \sigma_W$ and $\rho_{WZ} > \frac{(1+\alpha)\sigma_Z^2-(1-\alpha)\sigma_W^2}{2\alpha\sigma_W\sigma_Z}$,0 the variance in (12) is smaller than that in (7). As expected, the human involvement in the generation process of Y_t always results in an increase in the variability of ΔY_t, compared to the case in which the human intervention was completely absent.

(iv) It is easy to prove the following proposition (see Appendix): The case in which DM focuses on W instead of Z with a time-invariant reaction parameter α is equivalent to the case of a DM focusing on Z with a time-varying reaction parameter γ_t. This means that DM does not have to exhibit "deviant behaviour" in order to achieve her goals. The latter may be equivalently achieved if DM exhibits "politically correct" behaviour combined with time-varying degree of reaction.

4 Preferential Ambiguity with Learning

Let us now assume that the expert, utilizes the information that is being accumulated over time, to update his model by repeatedly by comparing the realized values of ΔY_t to those implied by his model. Retaining the assumption that the expert fails to observe the discrepancy between DM's and society's preferences, the only possible form of learning, is "parameter updating". More specifically, let the perceived (by the expert) law of motion (PLM) be

$$\Delta Y_t = AZ_{t-1} + u_t \tag{13}$$

where A is the parameter that the expert tries to estimate and u_t is a Gaussian IID process with zero mean. PLM is the reduced form model that the expert has in his mind when communicating his forecasts of ΔY_t, i.e. $E_{t-1}(\Delta Y_t)$, to the expert. To update the parameter A, he applies the recursive least squares (RLS) to (13). This methodology produces an estimate $\hat{A}_t$ for each time t, which minimizes the mean squared error, namely $E\left(\Delta Y_t - E_{t-1}(\Delta Y_t)\right)^2$ (For details see Appendix). It is easy to show that

$$\hat{A}_t \xrightarrow{p} A^* = -\frac{\alpha}{1-\alpha}\frac{\rho_{WZ}\sigma_W\sigma_Z + \mu_Z\mu_W}{\sigma_Z^2 + \mu_Z^2} \tag{14}$$

where "$\xrightarrow{p}$" signifies convergence-in-probability. Hence, asymptotically the expert's view on A will settle down on A^* (which is different than $-\frac{\alpha}{1-\alpha}$). Does the stochastic process ΔY_t converges-in-distribution? The answer is affirmative (see Appendix). Specifically, the asymptotic objective conditional distribution of ΔY_t is

$$\Delta Y_t \mid \mathcal{F}_{t-1} \sim N\left(a\left(A^* Z_{t-1} - W_{t-1}\right), \beta^2 \sigma_X^2\right), \tag{15}$$

whereas, the corresponding objective unconditional distribution is

$$\Delta Y_t \sim N\left(a\left(A^* \mu_Z - \mu_W\right), \left(\alpha A^*\right)^2 \sigma_Z^2 + \alpha^2 \sigma_W^2 + \beta^2 \sigma_X^2 - 2\alpha^2 A^* \rho_{WZ} \sigma_W \sigma_Z\right). \tag{16}$$

Comments

(i) The asymptotic parameter A^* can be decomposed in two terms: $-\frac{\alpha}{1-\alpha}$ and $\frac{\rho_{WZ}\sigma_W\sigma_Z+\mu_Z\mu_W}{\sigma_Z^2+\mu_Z^2}$. The second term may be thought of as an "adjustment factor" that captures the effects of learning. Specificallly, in the no-learning case the expert is always under the impression that the coefficient of Z_{t-1} is $-\frac{\alpha}{1-\alpha}$ (10) whereas under learning he ends up believing that this coefficient is A^*. As expected, when $W_t \equiv Z_t$, the second term is equal to 1 and, A^* collapses to $-\frac{\alpha}{1-\alpha}$, that is the coefficient of the benchmark case.

(ii) For each t, the objective distribution of ΔY_t does not coincide to the corresponding subjective distribution of the expert. More specifically, for each t, the difference between the objective conditional mean and the subjective conditional mean of ΔY_t is given by $(a-1) A_{t-1} Z_{t-1} - \alpha W_{t-1}$, which, in general, is different from 0. As a result, there exists a non-zero bias at each point t. The important question is whether this error asymptotically vanishes. The answer to this question is, in general, negative. The corresponding difference in the asymptotic means is $\alpha\left(\frac{\rho_{WZ}\sigma_W\sigma_Z+\mu_Z\mu_W}{\sigma_Z^2+\mu_Z^2}\mu_Z - \mu_W\right)$, which is zero iff $\rho_{WZ} = \frac{\mu_W\sigma_Z}{\mu_Z\sigma_W}$. This means that in spite of the learning process, the expert never achieves a full understanding of the situation, thus commiting a forecast error even asymptotically.

5 Comparison

The unconditional variance for all the above cases can be written as

$$\alpha^2\left[\left(\frac{\alpha}{1-\alpha}R\sigma_Z + \sigma_W\right)^2 - 2\frac{\alpha}{1-\alpha}R\left(1-\rho_{WZ}\right)\sigma_Z\sigma_W\right] + \beta^2\sigma_X^2$$

To arrive at the first case, we have that $R = \rho_{WZ} = 1, \sigma_W = \sigma_Z$. For the second case we need $R = 1$ and the third $R = \frac{\rho_{WZ}\sigma_W\sigma_Z+\mu_Z\mu_W}{\sigma_Z^2+\mu_Z^2}$. Note that the unconditional variance is increasing (decreasing) in R if $R > (<) - \frac{1-a}{\alpha}\frac{\sigma_W}{\sigma_Z}\rho_{WZ}$.

6 Introducing Deferential Ambiguity

Let us now introduce a second expert who (initially or permanently) disagrees with the first regarding the values of the structural parameters, θ, of the true model. More specifically, we assume that both experts know the true structural form of the model (including the DM's preference variable Z) but take different views on the values of its parameters, with none of these experts knowing the true value of θ. This case corresponds to the so-called "conflict ambiguity" in Baillon, Cabantous and Wakker (2012): "For the second source of uncertainty, each agent alone fully satisfies Bayesianism, with a precise probability judgment. However, the two agents give different judgments, generating ambiguity for the decision maker aggregating their beliefs. This source of uncertainty, which is characterized by between-agent ambiguity (heterogeneous beliefs), is called conflict (C-)ambiguity in this paper." (pp. 117). As already mentioned, this type of ambiguity does not affect exclusively DM; instead because of the endogeneity of DM's forecasts, conflict ambiguity produces a "boomerang effect" by injecting this ambiguity back into the process of forecast formation of experts. This is what we call "deferential ambiguity". Each expert does not know at each point in time whether DM will defer to him or his competitor. Hence each expert should account for this uncertainty by introducing it explicitly into his model for the generation of Y_t. How this can be achieved and what are its effects of this type of ambiguity on the distribution of ΔY_t are briefly analyzed below:

The assumptions that we make are the following: (i) DM adopts Z_{t-1} (rather than W_{t-1}) at t-1 (there is no preferential ambiguity). (ii) There are two experts who believe that the structural model is given by equations, (2), (3) and (4). The two experts agree on $\mu_Z, \sigma_Z^2, \sigma_X^2$ but disagree on α and β. Hence, the agent i's epistemic state is represented by $\theta_i = [\alpha_i, \beta_i, \mu_Z, \sigma_Z^2, \sigma_X^2]$, $i = 1, 2$. Without loss of generality, assume that $\alpha_2 > \alpha_1$. (iii) The objective probability that DM at t-1 defers to expert's i point forecast, $E_{t-1}^i(\Delta Y_t)$, is p_i. (iv) DM combines experts forecasts by means of a linear pool using p_1 and p_2 as weights. (v) Both experts know the objective deferential probabilities p_1 and p_2 as well as DM's aggregation rule.

Under the foregoing assumptions, we get

$$E_{t-1}^i(\Delta Y_t) = -\frac{\alpha_i}{1-\alpha_i} Z_{t-1}$$

Now the actual law of motion becomes

$$\Delta Y_t = -\alpha(p\frac{\alpha_1}{1-\alpha_1} + (1-p)\frac{\alpha_2}{1-\alpha_2} + 1)Z_{t-1} + \beta X_t.$$

The conditional distribution is

$$\Delta Y_t \mid \mathcal{F}_{t-1} \sim N\left(-\alpha\left(p\frac{\alpha_1}{1-\alpha_1}+(1-p)\frac{\alpha_2}{1-\alpha_2}+1\right)Z_{t-1}, \beta^2\sigma_X^2\right),$$

where $\mathcal{F}_t$ represents the information until t. The unconditional distribution is

$$\Delta Y_t \sim N\left(-\alpha\left(p\frac{\alpha_1}{1-\alpha_1}+(1-p)\frac{\alpha_2}{1-\alpha_2}+1\right)\mu_Z, \left(\alpha\left(p\frac{\alpha_1}{1-\alpha_1}+(1-p)\frac{\alpha_2}{1-\alpha_2}+1\right)\right)^2\sigma_Z^2+\beta^2\sigma_X^2\right).$$

Comments

(i) Since the coefficient $-\alpha$ is always positive, the unconditional mean of ΔY_t will always be positive (negative) whenever the mean, μ_Z, of Z is positive (negative). This means that if the society desires, on average, a positive (negative) change in next period's level of emissions, this desire will be translated (via DM's actions) into an actual average positive (change) change in emissions.

(ii) If $\alpha_1, \alpha_2 < \alpha$, the unconditional variance is always smaller than in the case of one expert and no deferential uncertainty. This case, corresponds to the case where both experts assume that the DM's actions will be more active than they actually are. As a result, their opinions will be closer to the "no action" case. Specifically, the term $\mathcal{E}_{t-1}(Y_t) - Y^*_{t-1} = \mathcal{E}_{t-1}(\Delta Y_t) - Z_{t-1}$ is closer to 0, i.e. the "no action" case. As a result, the DM will take less actions, which reduces the unconditional variance. If $\alpha_1 < \alpha < \alpha_2$, i.e. the first (second) expert assumes that the DM's actions will be more (less) significant, the unconditional variance will be larger when $p < \frac{(\alpha_2-\alpha)(1-\alpha_1)}{(\alpha_2-\alpha_1)(1-\alpha)}$. If $\alpha < \alpha_1 < \alpha_2$, then the unconditional variance in the case of two experts will always be larger.

7 Conclusions

In this paper, we approach the problem of ambiguity in climate change from a different angle than the one adopted in the recent relevant literature. The salient features of our approach are the following: (i) Ambiguity is an epistemic state which characterizes not only the decision maker but the scientific experts as well. We distinguish between preferential ambiguity, which is defined as the expert's uncertainty about DM's preference variables and deferential ambiguity, which arises in the case of multiple experts. Deferential ambiguity may be born by both DM and the experts and stems from the potential difficulty of DM to decide which of the experts should refer to. (ii) DM's ambiguity does not affect the formation of her prior probability function (which is the standard assumption in the ambiguity aversion literature). Instead, it affects the formation of DM's posterior distribution, in the sense that DM is uncertain about the piece of information that she should condition upon. As a result, DM's ambiguity is compatible with probabilistic sophistication. (iii) Both types of ambiguity have significant effects on the probabilistic properties of environmental policy variables. With respect to the policy relevant question of whether these types of ambiguity increase the probability of a "tail event", we show that the answer to this question depends on the probabilsitic properties of DM's preference variable compared to that of the society, on the extent to which DM learns from experience, on how DM combines experts' information and on the pattern of interaction between preferential and deferential ambiguity.

8 Appendix

Proof of Proposition: We want to find a process $\{\gamma_t\}$ such that $\alpha\left(E_{t-1}(\Delta Y_t) - W_{t-1}\right) + \beta X_t = \gamma_{t-1}\left(E_{t-1}(\Delta Y_t) - Z_{t-1}\right) + \beta X_t$. Solving for γ_{t-1}, we get

$$\gamma_{t-1} = \frac{\alpha\left(E_{t-1}(\Delta Y_t) - W_{t-1}\right)}{E_{t-1}(\Delta Y_t) - Z_{t-1}}$$

Define $\rho_t = \frac{W_t}{Z_t}$. Since $E_{t-1}(\Delta Y_t) = -\frac{\alpha}{1-\alpha} Z_{t-1}$ the above equation becomes

$$\gamma_{t-1} = \frac{\alpha\left(-\frac{\alpha}{1-\alpha} Z_{t-1} - \rho_{t-1} Z_{t-1}\right)}{-\frac{\alpha}{1-\alpha} Z_{t-1} - Z_{t-1}}$$

and as a result, $\gamma_t = a^2 + \rho_t\left(1-\alpha\right)$.

Recursive Least Squares: The least squares estimate is

$$A_t = \left(\sum_{s=1}^{t} Z_{s-1}^2\right)^{-1}\left(\sum_{s=1}^{t} Z_{s-1}\Delta Y_s\right)$$

More conveniently, the least squares estimates may be written in a recursive manner as

$$\begin{aligned} A_t &= A_{t-1} + t^{-1} R_t^{-1} Z_{t-1}(\Delta Y_t - A_{t-1} Z_{t-1}) \\ R_t &= R_{t-1} + t^{-1}(Z_{t-1}^2 - R_{t-1}) \end{aligned}$$

where $R_t = t^{-1}\left(\sum_{s=1}^{t} Z_{s-1}^2\right)$. The objective is to find the asymptotic value of A_t, denoted by A^*, and the conditions that lead to $A_t \to A^*$?

The expert's forecast of ΔY_t at time $t-1$, is given by $E_{t-1}(\Delta Y_t) = A_{t-1} Z_{t-1}$, which under (8) yields

$$\Delta Y_t = \alpha\left(A_{t-1} Z_{t-1} - W_{t-1}\right) + \beta X_t.$$

Hence, the RLS system can be written as

$$\begin{aligned}\phi_t &= \phi_{t-1} + t^{-1} R_t^{-1} Z_{t-1}((\alpha-1) A_{t-1} Z_{t-1} - \alpha W_{t-1} + \beta X_t) \\ R_t &= R_{t-1} + t^{-1}(z_{t-1} z'_{t-1} - R_{t-1})\end{aligned}$$

In order to apply the standard convergence results of stochastic recursive algorithms, we need to set $S_{t-1} = R_t$, in order for the term R_t^{-1} in the lhs of the first equation to be a lagged variable:

$$\begin{aligned}\phi_t &= \phi_{t-1} + t^{-1} S_{t-1}^{-1} Z_{t-1}((\alpha-1) A_{t-1} Z_{t-1} - \alpha W_{t-1} + \beta X_t) \\ S_t &= S_{t-1} + t^{-1}\left(\frac{t}{t+1}\right)(Z_t^2 - S_{t-1}).\end{aligned}$$

The associated ordinary differential equation (ODE) that governs stability of the system above is

$$\frac{d\Phi}{d\tau} = h(\Phi) = \lim_{t\to\infty} E\left(Q\left(t, \Phi, z_t\right)\right)$$

where $\Phi = (A, S)^{'}, z_t = (Z_{t-1}, W_{t-1}, X_t)$ and E denotes the expectation of $Q\left(t, \Phi, z_t\right)$ taken over the invariant distribution of z_t, for fixed $\Phi.Q\left(t, \Phi, z_t\right)$ is derived by the RLS system and is defined as

$$Q\left(t, \Phi, z_t\right) = \begin{pmatrix} S^{-1} Z_{t-1}((\alpha-1) A Z_{t-1} - \alpha W_{t-1} + \beta X_t) \\ \left(\frac{t}{t+1}\right)(Z_t^2 - S) \end{pmatrix}$$

It follows that

$$\begin{aligned}h_\theta(\Phi) &= \lim_{t\to\infty} E\left(S^{-1} Z_{t-1}((\alpha-1) A Z_{t-1} - \alpha W_{t-1} + \beta X_t)\right) \\ h_S(\Phi) &= \lim_{t\to\infty}\left(\frac{t}{t+1}\right) E(Z_t^2 - S) = \left(\sigma_Z^2 + \mu_Z^2\right) - S\end{aligned}$$

The second relationship gives $S \to \left(\sigma_Z^2 + \mu_Z^2\right)$, and therefore,

$$\begin{aligned}h_\theta(\Phi) &= \lim_{t\to\infty} E\left(\left(\sigma_Z^2 + \mu_Z^2\right)^{-1} Z_{t-1}((\alpha-1) A Z_{t-1} - \alpha W_{t-1} + \beta X_t)\right) = \\ &= (\alpha-1) A - \alpha\left(\sigma_Z^2 + \mu_Z^2\right)^{-1}\left(\rho_{WZ}\sigma_W\sigma_Z + \mu_Z\mu_W\right) \qquad (17)\end{aligned}$$

The ODE (17) gives the system

$$\dot{A} = (\alpha-1) A - \alpha\left(\sigma_Z^2 + \mu_Z^2\right)^{-1}\left(\rho_{WZ}\sigma_W\sigma_Z + \mu_Z\mu_W\right)$$

whose solution is given by

$$A^* = -\frac{\alpha}{1-\alpha}\frac{\rho_{WZ}\sigma_W\sigma_Z + \mu_Z\mu_W}{\sigma_Z^2 + \mu_Z^2}$$

The convergence to A^* is convergence in probability, i.e. $\forall\varepsilon > 0$, $\lim_{t\to\infty} P\left(|A_t - A^*| \geq \varepsilon\right) = 0$. The E-stability amounts to the condition $\alpha < 1$, which holds by assumption. Therefore, with probability 1, the system will converge to the equilibrium, irrespective of the initial estimations A_0.

References

Al-Najjar, N. and J. L. Weinstein, 2009, The Ambiguity Aversion Literature: A Critical Assessment, *Economics and Philosophy*, 25: 249-284.

Baillon A, Cabantous L, Wakker P.P., 2012, Aggregating imprecise or conflicting beliefs: An experimental investigation using modern ambiguity theories, *Journal of Risk and Uncertainty,* 44:115–147.

Carnap, R., 1950, Logical Foundations of Probability, Chicago, IL: University of Chicago Press.

Clemen, R.T. and R. L. Winkler, 1999, Combining probability distributions from experts in risk analysis, *Risk Analysis,* 19(2): 187–203.

Coenen, F., D. Huitema, and L. O'Toole, Jr., eds. 1998, Participation and the quality of environmental decision making, Dordrecht, Netherlands: Kluwer.

French, S., 1981, Updating of belief in the light of someone else's opinion, *Journal of the Royal Statistical Society,* Ser. A 143 43-48.

Gilboa, I. and D. Schmeidler, 1989, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics*, 18: 141-153.

Heal G. and A. Millner, 2014, Uncertainty and Decision Making in Climate Change Economics, *Review of Environmental Economics and Policy* 8.1, 120–137.

Jacobs R. A., 1995, Methods for combining experts' probability assessments, *Neural Computation*, vol. 7, pp. 867–888.

Klibanoff, P., M. Marinacci, and S. Mukerji, 2005, A Smooth Model of Decision Making under Ambiguity, *Econometrica*, 73: 1849-1892.

Lewis, D., 1980, A Subjectivist Guide to Objective Chance, In Richard C. Jeffrey, ed., *Studies in Inductive Logic and Probability*, vol. II. Berkeley: University of California Press. Reprinted with postscripts in Lewis 1986, pp. 83–132.

Lindley, D. V., 1985, Reconciliation of discrete probability distributions, *Bayesian Statistics* 2, North-Holland (Amsterdam).

McKitrick, Ross R., 2014, Climate Policy Implications of the Hiatus in Global Warming, Vancouver: Fraser Institute, October 2, 2014.

Meacham, C., 2007, Chance and the Dynamics of De Se Beliefs, Doctoral Dissertation.

Press, D., 1994, Democratic dilemmas in the age of ecology. Trees and toxics in the American West, Durham.

Schmeidler, D., 1989, Subjective Probability and Expected Utility without Additivity, *Econometrica*, 57: 571-587.

Tunney, R.J., and F.V. Ziegler, 2015, Surrogate utility estimation by long-term partners and unfamiliar dyads, *Frontiers in Psychology*, 6. doi: 10.3389/fpsyg.2015.00315 .