

Ambiguity and climate policy

Antony Millner^{*1,2}, Simon Dietz^{2,3}, and Geoffrey Heal^{4,5}

¹Department of Agricultural and Resource Economics, University of California, Berkeley

²Grantham Research Institute on Climate Change and the Environment, London School of Economics
and Political Science

³Department of Geography and Environment, London School of Economics and Political Science

⁴Columbia Business School, Columbia University

⁵National Bureau of Economic Research

December 2010

Abstract

Economic evaluation of climate policy traditionally treats uncertainty by appealing to expected utility theory. Yet our knowledge of the impacts of climate policy may not be of sufficient quality to justify probabilistic beliefs. In such circumstances, it has been argued that the axioms of expected utility theory may not be the correct standard of rationality. By contrast, several axiomatic frameworks have recently been proposed that account for ambiguous beliefs. In this paper, we apply static and dynamic versions of a smooth ambiguity model to climate mitigation policy. We obtain a general result on the comparative statics of optimal abatement and ambiguity aversion and illustrate this sufficient condition in some simple examples. We then extend our analysis to a more realistic, dynamic setting, and adapt a well-known empirical model of the climate-economy system to show that the value of emissions abatement increases as ambiguity aversion increases, and that this ‘ambiguity premium’ can in some plausible cases be very large.

1 Introduction

The literature on optimal climate change mitigation policy has thus far remained faithful to the long tradition of welfare analysis based on expected utility maximization. The integrated assessment models that are widely used for policy evaluation all have a common welfare-analytic core. Most studies employ deterministic models (Manne & Richels, 1992; Nordhaus, 2008; Tol, 1997), allowing for efficient determination of optimal policies, which are then subjected to sensitivity analysis in order to test their robustness to changes in model parameters. Other studies employ stochastic models (Hope, 2006), and generally do not find optimal policies, but rather provide welfare assessments of exogenously specified greenhouse gas emissions pathways. A few authors have

^{*}Email address for correspondence: a.millner@berkeley.edu. We thank Cameron Hepburn, Christian Traeger, Martin Weitzman, and Raphael Cael for very helpful comments, and Malte Meinshausen for supplying us with the empirical estimates of the climate sensitivity distributions.

combined these two approaches by solving stochastic-dynamic control problems to endogenously determine optimal policies that account for future risks (Pizer, 1999; Keller et al., 2004; Kelly & Kolstad, 1999). Thus, while there are several models of varying complexity and emphasis, they share a common commitment to the expected utility framework.

The reasons for the primacy of expected utility theory as a normative model of rational choice are well known to economists. Its axiomatic foundations have been developed by several authors (von Neumann & Morgenstern, 1944; Savage, 1954; Anscombe & Aumann, 1963). Savage’s presentation is widely considered the most satisfactory, since it derives both utility functions and subjective probabilities from primitive preferences over ‘acts’, i.e. maps between states and outcomes. Indeed his axioms are often considered to be synonymous with rational choice. Nevertheless, Savage himself took a cautious approach to his theory, suggesting that it should only be applied in small worlds, in which it is possible to ‘look before you leap’, i.e. imagine every possible contingency, and identify a complete ordering of acts over these contingencies (see Binmore, 2009, for a discussion). Further potential limitations on the domain of applicability of Savage’s theory were famously identified by Ellsberg (1961), who showed that when our state of knowledge is more accurately described as uncertainty rather than risk (we use these terms in the sense of Knight (1921)), we may wish to violate Savage’s second axiom, the ‘sure-thing principle’. A strong Bayesian would interpret Ellsberg’s results as a contribution to positive, rather than normative, decision theory. Bayesians believe that Savage’s axioms define rational choice, in which case the preferences Ellsberg observes in his investigations are deemed irrational, and all uncertainty is always describable by a unique probability distribution. This viewpoint has however been strongly contested. As noted by Ellsberg (1961) and Slovic & Tversky (1974), people often stick to choices that violate the sure-thing principle in Ellsberg’s choice experiments, even when this violation is pointed out to them. This is in stark contrast to other decision theoretic ‘paradoxes’, such as the Allais paradox, where people often revert to the prescriptions of expected utility once their violation of the axioms is explained. It has been argued, we think convincingly, that when our information about the world is incomplete, inconsistent, or nonexistent, Savage’s axioms need not be the correct standard of rationality (Gilboa et al., 2008, 2009), and it does not necessarily make sense to describe our state of knowledge with a unique probability distribution over states of the world¹.

Given these views, our assessment of the validity of the expected utility approach to the welfare analysis of climate change policy must depend on how structured our beliefs about the climate system² are. If we have sufficiently high quality information to justify probabilistic beliefs, then the approach adopted thus far in the literature is unequivocally useful. If not, we need to justify why this approach is a useful approximation, or attempt to define welfare measures that are true to our actual state of knowledge about the climate system, and reflect our preferences over bets with unknown probabilities. Very often a good way of justifying an approximation is to embed it in a more general framework, and show that the increased power of this framework does not materially alter the results achieved by the approximation. Thus, provided we suspect that our

¹Ellsberg himself emphasizes that ‘...either the postulates failed to be acceptable in those circumstances as normative rules, or they failed to predict reflective choices...But from either point of view, it would follow that *there would be simply no way to infer meaningful probabilities for those events from their choices*, and theories which purported to describe their uncertainty in terms of probabilities would be quite inapplicable in that area.’

²Our work focusses on uncertainty about a key parameter of the climate system – the climate sensitivity – and for the most part assumes that economic parameters are known. More on this later.

knowledge of the climate system is not very high quality, there would seem to be good reason to develop approaches to policy evaluation which account for uncertainty and not just risk, since these will either justify our reliance on existing methods, or provide appropriate tools for future work.

What is our state of knowledge about the climate system, and can it be described by unique probabilities? We feel it is important to break the state of scientific knowledge about climate into two categories: broad scientific principles, and detailed empirical predictions. In the first category belong concepts such as the laws of thermodynamics, fluid dynamics, and statements of fact such as ‘CO₂ traps outgoing long-wave radiation, causing warming’. We believe these principles to be unimpeachable. In the second category belong the sophisticated models scientists use to convert these principles into predictions – energy balance models (EBMs), earth systems models of intermediate complexity (EMICs), and full-scale general circulation models (GCMs). These models can be enormously complex, and attempt to predict, among other things, the response of the global climate to increases in the concentrations of greenhouse gases. Because of the complexity of their task, and the intrinsic difficulties of prediction in highly nonlinear multi-dimensional physical systems (see Smith (2002, 2007); Stainforth et al. (2007); Frame et al. (2007) for illuminating discussions of the scientific and philosophical challenges of climate prediction), these models are not always in agreement with one another. As an example of this, consider Figure 1, which plots the results of several recent studies’ attempts to use different models and observational data to estimate the climate sensitivity³, a critical parameter for estimating the response of the climate to increases in CO₂ concentrations which features prominently in integrated assessment models. From the figure it is clear that there are many inconsistent estimates of this important quantity. This suggests that we are indeed in an environment characterized by uncertainty, rather than risk. There are two standard responses to this assertion: Why not simply aggregate the different estimates into a single probability density (Bayesian response); and surely all the estimates are not equally valid, so why not simply choose the ‘best’ estimate (scientists’ response)? The point we are making is that the spread in the set of distributions is an indication of the intrinsic ambiguity in our state of knowledge. If decision makers perceive this ambiguity, and are averse to it as Ellsberg’s results suggest they will be, their beliefs *cannot* be described by a single probability distribution. How then should policy reflect this uncertainty?

Since Ellsberg’s work there has been a series of theoretical advances in decision theory which provide axiomatic representations of preferences that account for uncertainty or ambiguity. Seminal contributions include Arrow & Hurwicz (1977); Schmeidler (1989); Gilboa & Schmeidler (1989) and Klibanoff et al. (2005). These are elegant models, and have found application in several areas of economics, especially finance (e.g. Dow & da Costa Werlang, 1992; Bassett et al., 2004; Gollier, 2009; Bossaerts et al., 2010). Hansen & Sargent (2007) have applied similar techniques in macroeconomics in order to derive policies that are robust to model misspecification. Despite the appeal of these models in other areas of applied economics, they have only just begun to filter into the argument around climate policy. Henry & Henry (2002) is perhaps the first paper to view climate policy through this lens, and focusses on formalizing precautionary policies as policies that account for ambiguity in scientific knowledge. Lange & Treich (2008) provide some comparative

³The climate sensitivity is the amount by which global mean surface temperature rises for a doubling of CO₂ concentration, in equilibrium.




Figure 1: Estimated probability density functions for the climate sensitivity from a variety of published studies, collated by Meinshausen et al. (2009).

statics results on the effect of ambiguity on optimal abatement in a two-period model. Although not confined to climate applications, the work of Traeger (2009) and Gollier & Gierlinger (2008) on the effect of ambiguity aversion on the social discount rate is clearly relevant and has important implications for the assessment of mitigation investments.

In this paper we hope to provide a further step along the path sketched out by these authors.

In Section 2 we provide insight into the comparative statics of ambiguity aversion in stylized timeless models of greenhouse gas emissions abatement policy choice. Examining the static case helps to build intuition for the effect of ambiguity aversion on optimal decisions. We introduce a model of decision under ambiguity, analyze its basic properties, and derive a new and quite general condition that allows us to perform comparative statics. This condition is sufficient for an increase in ambiguity aversion to increase optimal abatement, or conversely, to decrease it. We then extend our comparative statics to two simple illustrative examples of abatement policy. The first example directly applies our sufficient condition to obtain conditions on the model's input assumption which ensure that optimal abatement is increasing in ambiguity aversion. Our sufficient condition is not satisfied in our second example, but it is possible to explicitly compute optimal abatement, and its relation with ambiguity aversion. In Section 3 we extend this analysis to the dynamic case, and attempt to understand how ambiguity aversion affects the welfare assessment of dynamic abatement policies. The dynamic aspects of ambiguity models are especially difficult, and we only consider their implications to a limited extent. In particular, we do not account for the dynamic resolution of ambiguity over time through learning, but rather focus on computing welfare measures

in two extreme cases – when ambiguity is expected to resolve completely after one time period, and when ambiguity persists unchanged for all time. These extremes allow us to bound the welfare effects of alternative abatement policies. We compute welfare measures for sample exogenous abatement pathways which are fed into the integrated assessment model DICE (Nordhaus, 2008) to generate consumption streams, and assess how ambiguity over the correct probability density for the climate sensitivity affects them. We show that the effect of ambiguity aversion on welfare evaluations depends sensitively on the damage function assumed in DICE, with steep damages at high temperatures giving rise to a large ‘ambiguity premium’ on abatement. We interpret these empirical results by appealing to theoretical work on the social discount rate under ambiguity. Section 4 discusses the results of our analysis, and concludes.

2 The smooth ambiguity model and optimal abatement

A potential difficulty with several of the decision models that account for ambiguity is that they do not achieve a separation between ambiguous beliefs and attitudes towards ambiguity. This was overcome by the contribution of Klibanoff et al. (2005), who provided a preference representation that separates tastes from beliefs, and allows us to parameterize attitudes to ambiguity via a differentiable function, in a manner analogous to the way utility functions represent risk preferences⁴. Their formalism is thus perfectly suited to understanding how different degrees of ambiguity aversion affect policies and welfare estimates. We introduce their model below.

Define an ‘act’ *à la* Savage as a map between states and outcomes (see e.g. Gilboa (2009) for a detailed explanation of Savage acts). Klibanoff et al. (2005) define a set of axioms for preferences over ambiguous acts, and show that when these axioms are satisfied act f is preferred to act g if and only if

$$\mathbf{E}_p\phi(\mathbf{E}_\pi u \circ f) > \mathbf{E}_p\phi(\mathbf{E}_\pi u \circ g), \quad (1)$$

where u is a von Neumann-Morgenstern utility function, ϕ is an increasing function, and p is a subjective second-order probability over a set Π of probability measures π that the decision maker (DM) deems to be relevant to her decision problem. When ϕ is concave, the DM can be said to be ambiguity averse, i.e. she dislikes mean-preserving spreads in the set of expected utilities implied by her model set Π . To immediately give this model a climatic interpretation, suppose that the choice variable is the level of abatement a of greenhouse gas emissions. Assume that there are m probability models of how abatement affects expected utility. For each model, write the expected utility obtained as a function of a as $EU_m(a)$. Now suppose that the DM is ambiguity averse, and that she has smooth ambiguity preferences *à la* Klibanoff et al. (2005). Let ϕ define her ambiguity preferences, i.e. $\phi' > 0$, $\phi'' < 0$, and p_m be the subjective second-order probability of model/prior m . Her objective function V is then:

$$V(a) := \sum_m p_m \phi(EU_m(a)). \quad (2)$$

⁴Ghirardato et al. (2004) achieve a similar separation between tastes and beliefs, but their framework is not well suited to the comparative statics analysis and dynamic applications we consider in this paper.

The first-order conditions can be written as

$$\sum_m \hat{p}_m(a^*) \left. \frac{dEU_m}{da} \right|_{a=a^*} = 0 \quad (3)$$

where a^* is the optimal abatement level, and we have defined the ambiguity-adjusted second-order probabilities $\hat{p}_m$ as:

$$\hat{p}_m(a^*) = \frac{\phi'(EU_m(a^*))p_m}{\sum_n \phi'(EU_n(a^*))p_n}. \quad (4)$$

Equation (3) just says that the weighted sum over models of marginal expected utility with respect to abatement should be zero, where the weighting factors are just the $\hat{p}_m$. It is identical to the condition one would obtain under ambiguity neutrality, except that p_m is replaced by $\hat{p}_m(a^*)$, the ambiguity-adjusted weight on model m .

If we look at (4), it is clear that the ambiguity weighting emphasizes those models that predict low expected utilities, since ϕ' is a decreasing function. Moreover, an increase in ambiguity aversion puts more weight on models with low expected utilities, and less on those with high expected utilities. This can be formalized using the concept of monotone likelihood ratios. Gollier & Gierlinger (2008) prove the following result:

Proposition 1. *Suppose there is a set $\mathcal{M}$ of possible models of cardinality M , and indexed by m . Without loss of generality, assume that the EU_m are ordered such that $EU_1 \leq \dots \leq EU_M$. Let $\phi_2 = f(\phi_1)$, where f is increasing and concave, and let $(\hat{p}_m^1)_{m \in \mathcal{M}}, (\hat{p}_m^2)_{m \in \mathcal{M}}$ be the ambiguity-adjusted second-order probabilities associated with ϕ_1 and ϕ_2 respectively, as given by equation (4). Then $(\hat{p}_m^1)_{m \in \mathcal{M}}$ dominates $(\hat{p}_m^2)_{m \in \mathcal{M}}$ in the sense of the monotone likelihood ratio order, i.e. $(\hat{p}_m^2/\hat{p}_m^1)_{m \in \mathcal{M}}$ is decreasing in m .*

The proof is very simple - the ratio $\hat{p}_m^2/\hat{p}_m^1 \propto f'(\phi_1(EU_m))$, where the proportionality constant is independent of m . Since ϕ is increasing, f' is decreasing, and the EU_m are increasing in m , the factor on the right decreases when m increases. This provides a simple characterization of the effect of increased ambiguity aversion on the first-order condition. It is important to stress however that the weights $(\hat{p}_m)_{m \in \mathcal{M}}$ are endogenous to the optimization problem, as they depend on a^* . Thus translating how these weights change into statements about how optimal abatement changes when ambiguity aversion increases is a non-trivial task in general.

General results on the comparative statics of ambiguity aversion are hard to come by, and depend on the properties of the sequence of functions $(EU_m(a))_{m \in \mathcal{M}}$. The following proposition defines conditions on these functions that allow us to ascertain the effect of an arbitrary increase in ambiguity aversion on optimal abatement:

Proposition 2. *Suppose that $\frac{d^2 EU_m}{da^2} < 0$ for all m , and assume that for every fixed value of a , the sequences $(EU_m(a))_{m \in \mathcal{M}}$ and $\left(\frac{dEU_m(a)}{da}\right)_{m \in \mathcal{M}}$ are anti-comonotonic⁵ (comonotonic) in m . Then an increase in ambiguity aversion increases (decreases) the optimal level of abatement for the objective (2).*

Proof. See Appendix A. □

⁵Two sequences are anti-comonotonic if one is increasing and the other is decreasing. They are comonotonic if they are both increasing, or both decreasing.




Figure 2: Examples of comonotonic (left) and anti-comonotonic (right) relationships between the sequences $(EU_m(a))_{m \in \mathcal{M}}$ and $\left(\frac{dEU_m(a)}{da}\right)_{m \in \mathcal{M}}$.

To understand the conditions on the functions $EU_m(a)$ in this proposition, consider the diagram in Figure 2, in which abatement a is plotted horizontally and $EU_m(a)$ is plotted vertically for a model set containing 4 distinct models. Anti-comonotonicity means that for each value of a the models with low expected utilities have high derivatives of expected utility with respect to a – this case is represented in the right panel of Figure 2. Hence when these conditions are satisfied an increase in abatement will reduce the spread of expected utilities across models and reduce the cost of ambiguity, so that a rise in ambiguity aversion leads to more abatement. Conversely, if the sequences are comonotomic then a decrease in abatement will decrease the spread between expected utilities (represented in the left panel of Figure 2) and so an increase in ambiguity aversion leads, rather counterintuitively, to a drop in abatement.

Note that an increase in ambiguity aversion always implies a policy change leading to a smaller spread of expected utilities. With anti-comonotonicity this means more abatement and with comonotonicity it means less. If abatement pays off most when consumption is low, it is likely that those models which have risk distributions that are skewed towards the low values of consumption (because, for example, they place a lot of weight on high values of climate sensitivity) have both low expected utility, and high marginal expected utility. Thus the relationship between expected utility and marginal expected utility is likely to be anti-comonotonic in this case. Conversely, if abatement pays off least when consumption is low, marginal expected utility will be low in those models which place a lot of weight on low consumption, i.e. when expected utility is low; we are now in the realm of comonotonicity. Which (if any) condition holds in practice is an empirical question. Proposition 2 requires all the graphs of $EU_m(a)$ to be non-intersecting, and for the distance between any consecutive pair of expected utilities to be decreasing (or increasing in the case of comonotonicity) in a . It is quite possible that neither condition is satisfied, so we next present two examples to investigate further the comparative statics of ambiguity aversion. The

first example satisfies the conditions of proposition 2, while the second does not, but can be solved analytically to yield some insight into the comparative statics. In this latter example we will see that it is once again the effect of abatement on the spread in expected utilities that determines the comparative statics of ambiguity aversion.

Example 1. Suppose that we have two models of the effect of greenhouse gases on climate:

- Model 1: Climate is not sensitive to anthropogenic emissions.
- Model 2: Climate is sensitive to anthropogenic emissions. Suppose that, conditional on this model being correct, there are two climate states, a high damage state H , and a low damage state L . Abatement increases the probability of state L materializing.

The DM attaches a default utility level u_0 to today's climate. Let $u_0 - d_H$ be the utility level in the high climate damages state, and $u_0 - d_L$ be the utility level in the low damages state, with $d_H > d_L$. The DM must decide the level of abatement, a . The utility cost of abatement is given by a function $\Lambda(a)$, and let $\pi(a)$ be the probability of the low-damages state occurring as a function of abatement. Clearly we need $\pi'(a) > 0$, $\Lambda'(a) > 0$. Proposition 2 allows us to obtain the following:

Proposition 3. *Suppose that $\pi''(a) < 0$ and $\Lambda''(a) > 0$. Then an increase in ambiguity aversion increases optimal abatement in Example 1.*

Proof. The expected utilities obtained in models 1 and 2 are:

$$EU_1 = u_0 - \Lambda(a) \tag{5}$$

$$\begin{aligned} EU_2 &= \pi(a)(u_0 - d_L - \Lambda(a)) + (1 - \pi(a))(u_0 - d_H - \Lambda(a)) \\ &= u_0 - \Lambda(a) + \pi(a)(d_H - d_L) - d_H. \end{aligned} \tag{6}$$

Clearly, $EU_1 > EU_2$, since $\pi(a) \in [0, 1]$. In addition, simple differentiation shows that $\frac{dEU_1}{da} < \frac{dEU_2}{da}$. Thus the sequences (EU_1, EU_2) and $(\frac{dEU_1}{da}, \frac{dEU_2}{da})$ are anti-comonotonic for all a . Finally, the conditions $\pi''(a) < 0$, $\Lambda''(a) > 0$ ensure that $\frac{d^2 EU_m}{da^2} < 0$, $m = 1, 2$. Thus the conditions of proposition 2 are satisfied, and the result is established. $\square$

Example 2. Consider a stylized abatement problem, an extension of an example examined by Gollier (2009). Assume we have a set of risk distributions $\pi_\theta(x)$, indexed by the continuous parameter θ , where x is the magnitude of additive consumption damages from climate change. We can moderate these risks by investing in abatement a . Abatement reduces the magnitude of damages by a multiplicative factor that is linear in a . The effectiveness of abatement is controlled by a constant b_1 , so that damages averted at abatement level a are given by $b_1 a x$. Abatement comes at a consumption cost that is linear in a , with b_2 the cost per unit of abatement. We are ambiguity averse, so the objective function is:

$$V(a) = \int p(\theta) \phi \left(\int U(c_0 + (1 - b_1 a)x - b_2 a) \pi_\theta(x) dx \right) d\theta. \tag{7}$$

Proposition 4. *Assume that*

$$U(c) = -\frac{1}{A} \exp(-Ac) \quad (8)$$

$$\phi(U) = -(-U)^{1+\xi}/(1+\xi) \quad (9)$$

$$x \sim \mathcal{N}(\theta, \sigma) \quad (10)$$

$$\theta \sim \mathcal{N}(\mu, \sigma_0) \quad (11)$$

where $A, \xi, \sigma, \sigma_0, b_1, b_2$ and c_0 are non-negative constants, and the sign of μ is arbitrary. Then when $-b_1\mu/b_2$ is greater (less) than 1 the abatement level a^* that maximizes (7) is increasing (decreasing) in ξ . Moreover, a change in ξ only affects the mean of the ambiguity-adjusted second-order probability distribution $\hat{p}(\theta)$, which is decreasing (increasing) in ξ when $-b_1\mu/b_2$ is greater (less) than 1.

Before proving the proposition we unpack the interpretation of the parameters: our choices for U and ϕ correspond to constant absolute risk aversion (parametrized by A) and constant relative ambiguity aversion (parameterized by ξ) respectively. We have assumed that each of the risks $\pi_\theta(x)$ is a normal distribution, with fixed standard deviation σ , and an uncertain mean θ . It can be shown that under these assumptions the expected utilities $(EU_\theta(a))_{\theta \in \mathbb{R}}$ are concave in a , but do not satisfy either of the comonotonicity properties required by proposition 2. The second-order distribution $p(\theta)$ of the means θ is also normal, with mean μ and standard deviation σ_0 . Note that since x is intended to represent negative climate impacts, we implicitly assume that $\mu < 0$, although this condition is not required in any of our results. Even with $\mu < 0$ however, for each risk distribution there is a non-zero probability that impacts will be positive, since the normal distribution has support on the whole real line – this represents the possibility of climate windfalls. Finally, b_1 determines the effectiveness of abatement (marginal benefits of abatement are just $-b_1x$), b_2 is the marginal cost of abatement, and c_0 is a baseline consumption level. We now turn to a proof of the proposition:

Proof. $V(a)$ can be computed explicitly (see Appendix B), and maximized to show that the optimal value of a is

$$a^* = \frac{1}{b_1} + \frac{b_1\mu + b_2}{Ab_1(\sigma^2 + (1+\xi)\sigma_0^2)}. \quad (12)$$

Thus when $-b_1\mu/b_2$ is greater (less) than 1, the optimal value a^* is increasing (decreasing) in ξ . Some standard computations (see Appendix B) then show that

$$\begin{aligned} \hat{p}(\theta) &\sim \mathcal{N}(\mu - \sigma_0^2\xi A(1 - b_1a^*), \sigma_0^2) \\ &= \mathcal{N}\left(\mu + \frac{\sigma_0^2\xi(b_1\mu + b_2)}{\sigma^2 + (1+\xi)\sigma_0^2}, \sigma_0^2\right) \end{aligned} \quad (13)$$

Thus in this case, the only effect of ambiguity aversion is to shift the mean of the second-order probabilities $\hat{p}(\theta)$ downwards (upwards) when $-b_1\mu/b_2$ is greater (less) than 1. Thus when $-b_1\mu/b_2 > 1$, as ambiguity aversion increases, increased weight is placed on those risk distributions with mean $\theta < \mu$, while the weights on risks with $\theta > \mu$ decline. The converse obtains when $-b_1\mu/b_2 < 1$. $\square$

The condition $-b_1\mu/b_2 > 1$ in the proposition is intuitively compelling. Recall that the benefit

of abatement a when damages are x is just $-b_1xa$, and the costs are b_2a . Thus we may read our condition as (Expected benefit)/Cost > 1 . The result says that if the expected benefit of abatement outweighs costs, then ambiguity aversion increases optimal abatement. An increase in ambiguity aversion means that the DM is more averse to spreads in expected utility, so in order to understand the comparative statics we need to understand how a change in a affects the spread in expected utilities. If, on average, an increase in a increases consumption (i.e the argument of U in (7)), the spread in expected utilities will decrease, since by concavity U , and therefore its expected value, becomes less sensitive to variations in consumption when consumption increases. For a fixed value of x , the argument of U increases with a if and only if $-b_1x/b_2 > 1$. Thus our condition $-b_1\mu/b_2 > 1$ may be interpreted as ensuring that on average an increase in a increases consumption, which in turn reduces the spread in expected utilities. Since a reduction of the spread in expected utilities is desirable when ξ increases, it is optimal to increase a . When $-b_1\mu/b_2 < 1$ the converse case obtains, with a decrease in a increasing consumption on average, and thus being desirable when ξ increases. The fact that in this case the effect of ambiguity aversion depends only on the benefit/cost ratio of abatement is due to the linearity of consumption in a .

3 Evaluating dynamic abatement pathways under ambiguity

The preceding section abstracted the climate change abatement problem to a high level. Perhaps most importantly, the models examined thus far have all been atemporal. While static models are useful for gaining an intuition for the new effects that ambiguity aversion introduces into familiar problems, they are of limited use for deriving results that are meaningful for climate policy. The abatement problem is in its essentials a dynamic decision problem, in that it requires us to trade off near-term costs against uncertain long-term benefits. This section examines how ambiguity over the future benefits of abatement affects the welfare analysis of alternative climate policies. We first introduce a dynamic extension of the smooth ambiguity model of Klibanoff et al. (2005), derived in Klibanoff et al. (2009). Doing so immediately raises issues of how dynamic consistency and learning are incorporated in such a model, and we discuss these. Second, we make an empirical application of the Klibanoff et al. (2009) model to climate policy, using the DICE integrated assessment model (Nordhaus, 2008) to generate consumption streams given choices of abatement strategy, in the face of ambiguous information about the climatic response to greenhouse gas emissions (i.e. represented by the climate sensitivity parameter). We proceed directly to these numerical simulations, since in the full dynamic context there is a limited amount that can be deduced about the relationship between abatement and ambiguity aversion analytically: much depends on the empirical details.

3.1 Dynamic welfare functions

Klibanoff et al. (2009) obtain a representation of preferences over time- and state-dependent acts, i.e. contingent plans that map the nodes of a decision tree into consumption streams. If we let $s^t = (x_1, \dots, x_t) \in \Gamma$ denote a decision node, where $x_\tau \in \mathcal{X}_\tau$ is an observation at time τ and Γ is the set of all nodes, then a generic plan f maps s^t into consumption. Klibanoff et al. (2009) show

that preferences over plans that satisfy consequentiality, dynamic consistency, and further axioms that are similar to those employed in the static representation result, can be represented by

$$V_{s_t}(f) = u(f(s^t)) + \beta \phi^{-1} \left[\int_{\Theta} \phi \left(\int_{\mathcal{X}_{t+1}} V_{(s^t, x_{t+1})}(f) d\pi_{\theta}(x_{t+1}; s^t) \right) dp(\theta|s^t) \right], \quad (14)$$

where $\theta \in \Theta$ indexes the set of alternative probability models, and β is a discount factor. They refer to this representation as the *recursive form*, for obvious reasons. Although the focus of their paper is on achieving this result, Klibanoff et al. (2009) also derive an alternative, distinct, preference representation, which they call the *reduced form*:

$$\hat{V}_{s_t}(f) = u(f(s^t)) + \beta \phi^{-1} \left[\int_{\Theta} \phi \left(\sum_{s \in \Gamma \cap \{s^t\}} \pi_{\theta}(s|s^t) \sum_{\tau \geq t+1} \beta^{\tau-(t+1)} u(f(s)(\tau)) \right) dp(\theta|s^t) \right]. \quad (15)$$

The difference between these two representations arises from the different requirements they impose on consistency between so-called first- and second-order acts (see Klibanoff et al., 2009, for details). The recursive form requires consistency only for ‘one-step-ahead continuation plans’, while the reduced form requires consistency for all continuation plans. It turns out that the recursive form is dynamically consistent, while the reduced form is not. To see why notice that the reduced form evaluates all future consumption streams by averaging over the *current* second-order uncertainty distribution $p(\theta|s^t)$. However it should be clear that this is not the distribution that we would use to compute expectations at a future decision node, since in reaching that future node new observations would be made that allow us to update the second-order probabilities – this constitutes a violation of dynamic consistency. In fact, the reduced form representation requires DMs to act as if all the second-order uncertainty will be resolved in the next time period. The recursive formulation respects dynamic consistency, since it only uses the current information to average over continuation values, which are themselves recursively defined in terms of averages over future (updated) information. It is thus capable of representing preferences that coherently account for the persistence of ambiguity. The conflict between these two distinct representations is a new feature of inter-temporal choice which arises from the desire to simultaneously represent ambiguity-sensitive preferences and respect consequentialism (Klibanoff et al., 2009, p. 946). Note that when ϕ is affine, i.e. the DM is ambiguity neutral, the two representations are identical; this explains why these issues are not encountered in the standard risk case.

While the recursive form (14) can be treated as a Bellman-like equation that can be used to solve for optimal abatement policies once a noise distribution for temperature observations has been specified, this approach entails considerable computational complications (see e.g. Kelly & Kolstad, 1999). This is due to the high dimensionality of the state space when there are multiple distributions for the climate sensitivity that need to be tracked and updated as observations of temperature are made⁶. We hope to take up this task in future work. Instead of solving for

⁶To get an idea of what a full dynamic analysis might entail, suppose that there are 20 possible distributions for climate sensitivity, each of which comes with a second order weight (19 independent weights in total), and that each distribution can be described by two sufficient statistics, each of which is updated with each new observation. Then there are $2 \times 20 + 19 = 59$ ‘informational’ state variables which appear as arguments of the value function. The curse of dimensionality means that obtaining solutions to dynamic programming problems of this size is not possible by conventional computational methods.

optimal policies, we will content ourselves with evaluating exogenous policies, and assessing the dependence of these evaluations on ambiguity. Evaluating exogenous policies is a desirable step – emissions pathways may, for example, be subject to political constraints that prevent them from being ‘optimal’, yet we may still wish to evaluate their effects on welfare, and determine how those effects depend on our attitude to ambiguity.

It will be convenient in what follows to have a means of representing the welfare difference between two policies in consumption units, since this makes welfare changes for different preferences directly comparable. Because we must potentially deal with non-marginal changes in welfare (because very high climate sensitivities could lead to very large climate damages), the measure we will use in order to convert welfare changes into consumption units is the Stationary Equivalent (SE) (Weitzman, 1976). The SE of a welfare function V is defined as the value of consumption $c(V)$ which, when held constant, gives rise to welfare equivalent to V . Thus we define $c(V)$ through

$$\sum_{t=0} \beta^t P_t u(c(V)) = V, \quad (16)$$

where P_t is the population at time t . Define the fractional change in the SE induced by an abatement policy, relative to a business as usual (hereafter BAU) baseline, as Δ . A simple calculation shows that when u is of the constant relative risk aversion (CRRA) form (see (19) below), Δ is given by:

$$\Delta := \frac{c(V_{ABATE}) - c(V_{BAU})}{c(V_{BAU})} = \left(\frac{V_{ABATE}}{V_{BAU}} \right)^{\frac{1}{1-\eta}} - 1, \quad (17)$$

where $ABATE$ denotes a generic abatement policy. All the results that follow use Δ to represent welfare differences.

3.2 The DICE model, policies, and preferences.

To make the discussion concrete, we will now use the DICE model to analyse the empirical effect of ambiguity on climate policy, focusing on ambiguity over the climate sensitivity. We will not be using the DICE model to solve for optimal abatement policies, but will rather use it to generate consumption streams that depend on exogenous policy settings, and a value for the climate sensitivity.

If we denote the climate sensitivity by S , the exogenous savings rate by $\sigma(t)$, and the exogenous abatement effort by $a(t)$, then we view DICE as the following function:

$$\text{DICE}(S; \sigma(t), a(t)) = c(t), \quad (18)$$

where $c(t)$ is a consumption stream. We can now compute this function for a variety of values of S , holding $\sigma(t)$ and $a(t)$ constant, to see how the consumption stream depends on the climate sensitivity. Before specifying our choice of the controls, a few words are in order about the DICE model, which is explained in full detail in Nordhaus (2008). DICE is a well known integrated assessment model of the connections between economic activity and climate change. A standard Ramsey-Cass-Koopmans growth model with aggregate capital and labour inputs is linked to climate change through emissions of greenhouse gases, which cause global warming and, with a lag, reduce output by means of a reduced-form ‘damage function’ (more on this later). This damage

function incorporates assumptions about adaptation to climate change, which can reduce output losses, leaving the representative agent with the choice of how much to invest in emissions abatement $a(t)$, as well, of course, as how much to invest in the composite capital good. The model includes many economic and climate parameters, all of which are at least to some extent uncertain. However, for the sake of tractability, we focus on uncertainty (specifically ambiguity) about the climate sensitivity S . This is justified, since S is known to be one of the most important uncertainty parameters in integrated assessment of climate change, with a potentially strong influence on the value of abatement strategies (Weitzman, 2009). Moreover, as we have argued above, it is a parameter about which our knowledge is appropriately represented as ambiguous⁷. Future work could also, where appropriate, treat knowledge about other parameters as ambiguous.

In our empirical results below, we pick specific values for the controls $\sigma(t)$ and $a(t)$. For simplicity, we assume that the savings rate $\sigma(t)$ is a constant 22%. In order to fix the abatement effort $a(t)$, we consider three emissions abatement scenarios. Abatement effort is represented in DICE by the emissions control rate, a number between 0 and 1, which controls the emissions intensity of gross economic output (i.e. before climate damages are incurred). When the control rate is $a(t)$, a fraction $1 - a(t)$ of gross output contributes to emissions. Our three scenarios for the control rate are a ‘Business as usual’ scenario, a scenario that limits the atmospheric concentration of CO₂ to twice its pre-industrial level (560 parts per million, hereafter referred to as the 2 CO₂ scenario), and a more aggressive abatement scenario that limits the concentration of CO₂ to only one-and-a-half times its pre-industrial level (420ppm, hereafter referred to as the 1.5 CO₂ scenario)⁸. Both of the abatement scenarios have been prominent in recent international negotiations about climate policy. The emissions control rates corresponding to our three scenarios are depicted in Figure 3. They are each taken from Nordhaus (2008)⁹.

Suppose now that we have a set of M distributions $\{\pi_m(S)\}$ for the climate sensitivity S , indexed by m , with initial second-order probability weights given by p_m . How can we compare the welfare effects of the two abatement pathways? Ideally we would like to use the recursive preference representation (14), however we immediately run into difficulties. In order to apply this representation we need to know how the first- and second-order probabilities $(\pi_m(S), p_m)$ change over time as more information is revealed. Yet this information is not part of our primitive inputs to the model – all we know is the values of our plan $f(s^t)$, but not how the sequence of realized states $\{x_t\}_{t=1\ldots\infty}$ affects beliefs $\pi_m(S)$ and p_m . Moreover, strictly speaking, the manner in which we make use of the DICE model implies that *all* uncertainty (i.e. risk and ambiguity) is

⁷In fact, there are good reasons to believe that the upper tails of the various distributions of S estimated in the climate science literature are not well constrained by observations, and depend heavily on scientists’ subjective choice of prior (Allen et al., 2006). We are arguing that scientists’ prior beliefs about S are not sufficiently structured for them to be representable by a unique probability distribution; hence unique posteriors also do not exist.

⁸Note that the control rates associated with the two abatement scenarios are designed so that they achieve the specified stabilization targets in an idealized run of DICE in which damages are zero for concentrations below the stabilization target, and rise sharply to very high values above the target. This is the method used by Nordhaus (2008) to generate controls that achieve a given stabilization target, however it should be born in mind that these controls will not in general achieve this target when they are used as inputs to DICE for alternative damage functions, or values of S . Since all we require for our purposes is a plausible choice of controls, we need not concern ourselves too much with whether they achieve a given stabilization target.

⁹It may be seen that there are kinks in the emissions control rates corresponding to the 2 CO₂ and 1.5 CO₂ scenarios, which are taken directly from Nordhaus (2008). The kink in 2015 that is especially evident in the 2 CO₂ scenario is due to the fact that the control rate in 2005 is not allowed to vary. Nordhaus does not explain the other kinks that are evident. He also only states the control values until 2105. We have linearly extrapolated the controls to 2215, assuming a control rate of 1 at 2215 for the abatement policies.




Figure 3: Emissions control rates for our sample abatement scenarios.

resolved immediately. This is so since in our setup the DICE model is a deterministic function. This implies that as soon as a single time period has elapsed the DM can use the known model equations to work back from the changes in the state variables to deduce the exact value of the climate sensitivity¹⁰. Clearly this does not map onto the scientific situation in reality, in which uncertainty about the climate sensitivity is likely to persist for some time.

There are two resolutions to these difficulties. The first is to use the reduced-form representation to evaluate welfare. As mentioned above, this implies that we are evaluating policies today as if all ambiguity were to be resolved in the very next time step. This has the disadvantage of not respecting dynamic consistency (except of course in the unlikely event that we believe ambiguity really will be resolved after one time step), but has the advantage of being consistent with our usage of the DICE model as a deterministic function. The reduced form approximation can be made more realistic by increasing the length of the time step, so that an *a priori* plausible amount of time passes before ambiguity is resolved. Thus in our computations below we increase the length of the time step in DICE from one to three decades. Since DICE's initial period is 2005, ambiguity is assumed to be resolved in 2035 for the reduced form welfare measure.

The second resolution is to use the recursive form, and assume an *ad hoc* model for how beliefs about ambiguity and risks are updated over time. This has the advantage of allowing for a persistent (albeit arbitrary) dynamic representation of ambiguity, but requires us to pretend that the consumption streams at our disposal are divorced from the model that generated them, i.e. we pretend that we do not know the model equations, so that uncertainty is not resolved in the first time step. An extreme version of this strategy would be to assume that no information is

¹⁰This is so provided the model's state equations are 'invertible' (i.e. we can infer parameter values from sequential realizations of the state variables). The immediate resolution of uncertainty in our model is due to the fact that there are no noise terms in the state equations.

obtained at any point in time, i.e. ambiguity is perfectly persistent, and the DM never updates her beliefs. In addition, we need to assume that the DM believes that the system will be in the same state at every time step, i.e. the probability tree she faces at every decision node is identical. These are not terribly realistic assumptions, as they require us to make welfare evaluations as if the DM cannot anticipate the effect of the dynamics on the system's states and her beliefs. They will however be sufficient to capture some of the effects of persistent ambiguity on the welfare evaluation.

Our strategy will be to use both the reduced form and the recursive form with perfectly persistent ambiguity to compute welfare measures. It seems reasonable to believe that these two approaches provide meaningful bounds on the value of an abatement project. Any model which admits updating of probabilities must fall somewhere between these two extremes¹¹.

For the sake of argument we assume equal second-order weights on each of the distributions for climate sensitivity depicted in Figure 1, and choose u and ϕ both to be isoelastic functions (constant relative risk and ambiguity aversion respectively):

$$u(c) = \frac{c^{1-\eta}}{1-\eta}, \quad (19)$$

$$\phi(u) = \begin{cases} \frac{U^{1-\xi}}{1-\xi} & \eta < 1 \\ \frac{-(-U)^{1+\xi}}{1+\xi} & \eta > 1 \end{cases}, \quad (20)$$

where η (ξ) is the coefficient of relative risk (ambiguity) aversion. Note that since ϕ operates on utility, which is measured in different units to consumption and will usually be estimated over a quite different range of absolute values, larger values of ξ are considered plausible, compared with η . For instance, Gollier & Gierlinger (2008) calibrated the above model by considering a simple binary lottery with unknown probabilities. A uniform distribution over the second-order probabilities of the two outcomes is assumed, and the ambiguity premium is set at 10% of the expected value of the lottery, based on the experimental results reported in Camerer (1999). Calibrating the model at $\eta = 2$, they found $\xi \in [5, 10]$ to be roughly plausible. This is roughly consistent with the analysis of asset returns under ambiguity in Ju & Miao (2009), which found $\xi = 8.864$. We will take the upper estimate of $\xi = 10$ as our representative case of an ambiguity-averse decision maker, however we note that in general the issue of calibrating the ambiguity model is a difficult one, and the method described above is by no means the last word on the subject¹². Consequently in some of our figures below we report a wider range, $\xi \in [0, 50]$.

We then use DICE to generate consumption streams for climate sensitivities in the range $[1, 10]$ ¹³, and compute welfare measures for each of the two representations of ambiguity and abatement strategies. In all our computations, we set the annual discount rate on utility to 0.1%, consistent with recent viewpoints by, for example, Stern (2008), Dasgupta (2008), and Heal (2009).

¹¹The fact that the recursive form is equivalent to the reduced form when the learning process employed in the recursive model resolves all uncertainty in one time step is a trivial consequence of the fact that the two representations agree on their ranking of one-step-ahead continuation plans.

¹²To understand some of the complexity of the calibration problem for the ambiguity model, note that since the representation (1) is not linear in the utilities u in general, the calibration will depend on the choice of the normalization of utility.

¹³The simple climate model built into DICE breaks down for $S < 1$, so we are forced to truncate the distributions in Figure 1, applying a lower bound value of 1. We do not lose much information in this way, since inspection of Figure 1 shows that there is a very low probability that $S < 1$, and that the upper tail is of much greater importance.

In our base case, we set the coefficient of relative risk aversion (η) to 2, and we use the standard DICE damage function,

$$\Omega(t) = \frac{1}{1 + \alpha_1 T(t) + \alpha_2 T(t)^2}, \quad (21)$$

where T is the increase in global mean temperature above the pre-industrial level, and α_1 and α_2 are coefficients. The damage from warming as a fraction of GDP is $1 - \Omega(t)$, which with the default calibration of $\alpha_1=0$ and $\alpha_2=0.0028$, gives damages of 1.7% of GDP for 2.5°C warming and 6.5% for 5°C warming (we comment on the plausibility of these damages below).

Figure 4 plots Δ for the 1.5 CO₂ and 2 CO₂ abatement strategies as a function of the coefficient of relative ambiguity aversion (ξ) for our base case. The figure shows that Δ increases as ξ increases. Focussing on the recursive form with perfectly persistent ambiguity (blue), the Δ of the 1.5 CO₂ strategy increases from approximately 0.09% when $\xi=0$ to approximately 0.25% when $\xi=50$. Under the 2 CO₂ strategy recursive Δ increases from 0.63% when $\xi=0$ to 0.76% when $\xi=50$.

That increasing ambiguity aversion has this effect on Δ essentially reflects the forces at play in our comparative statics analysis above. Climate is sensitive to emissions, so the higher are emissions, the more likely it is that high damages result. By proposition 1 an increase in ambiguity aversion places more weight on models with lower expected utility, or in other words models in which damages are higher. In these models, costly emissions abatement is of more value as it avoids greater amounts of climate damage. In Figure 4 the base-case relationships appear close to linear for the range of ξ that we consider. However it is easy to see that, were it computationally feasible¹⁴ to increase ξ indefinitely, the SE of the policies, and hence Δ , would eventually asymptote to a constant value¹⁵.

Figure 5 considers the relative importance of ambiguity aversion compared with risk aversion, and how the two preferences interact with one another. It plots Δ as a function of η for $\xi = 0$ and $\xi = 10$. We discuss the interpretation of this figure in the following sub-section.

3.3 Heuristic explanation of results, and the role of the damage function

In order to understand the qualitative features of Figure 5 we contrast it with the expression for the certainty-equivalent social discount rate (ρ) under ambiguity aversion, derived in Traeger (2009) and Gollier & Gierlinger (2008). These derivations assume isoelastic forms for u and ϕ , and that consumption grows at an uncertain rate $g \sim \mathcal{N}(\theta, \sigma)$, where the mean growth rate θ is itself uncertain and distributed according to a second-order probability distribution, $\theta \sim \mathcal{N}(\mu, \sigma_0)$. It can then be shown that

$$\rho = \delta + \eta\mu - \frac{\eta^2}{2}(\sigma^2 + \sigma_0^2) - \xi|1 - \eta^2|\frac{\sigma_0^2}{2}, \quad (22)$$

where δ is the utility discount rate. The first two terms in this expression are familiar from the standard Ramsey formula for the social discount rate under certainty, and capture inter-temporal substitution effects. The third term is the standard correction due to the uncertainty in the growth

¹⁴The values of ξ it is feasible to consider numerically are limited by the accuracy of floating point arithmetic.

¹⁵The limiting values as $\xi \rightarrow \infty$ are not the same for the recursive and reduced form welfare measures. The limiting reduced form SE depends only on the model with the lowest discounted expected utility, while the limiting recursive form SE depends on the discounted sum of the lowest expected utility of *any* model at each point in time.




Figure 4: Percentage change in SE of abatement relative to BAU for 1.5 CO₂ (420ppm) and 2 CO₂ (560ppm) abatement pathways, as a function of the coefficient of relative ambiguity aversion (ξ). Estimates based on the reduced form model are plotted in red, while estimates based on the recursive model with perfectly persistent ambiguity are plotted in blue. $\eta = 2, \delta = 0.1\%$ for these simulations.

rate (note that the variance of the growth rate is the variance of the composite distribution that arises from the combination of uncertainty about g and θ). The final term is a new addition due to ambiguity aversion.

Since differences in the expected utilities of abatement and BAU pathways only manifest significantly in the distant future (when damages are large), it is intuitive that the social discount rate may play a useful role in explaining our results. If ρ is large, we expect welfare differences between abatement and BAU policies to be small. Similarly, those values of η for which ρ is highly sensitive to ξ are also likely to be the values of η for which we see a significant effect of ambiguity aversion on the welfare difference measure Δ . Although the assumptions about the first- and second-order uncertainty on consumption growth rates used to derive (22) do not map neatly onto our empirical application, the expression nevertheless provides useful qualitative insights.

First notice that ρ is decreasing in ξ – this is consistent with our finding in Figure 4 that Δ is increasing in ξ . Second, consider the case $\eta = 0$, for which $\rho = \delta - \frac{1}{2}\xi\sigma_0^2$. It is clear that in this case an increase in ξ can potentially have a large impact on the valuation – in fact if ξ is large enough ρ can be negative, thus placing increasing weight on the future, which should lead to large estimates of NPV. This is reflected in Figure 5, where we see that the difference between $\Delta_{\xi=0}$ and $\Delta_{\xi=10}$ is at its largest when $\eta = 0$. Third, consider the case $\eta = 1$. It is clear that in this case ρ is independent of ξ , implying that ambiguity aversion has no effect on welfare calculations. This is also reflected in Figure 5, where we see that $\Delta_{\xi=0}$ and $\Delta_{\xi=10}$ coincide at $\eta = 1$. Fourth, notice that when η becomes large, Δ tends to zero for both $\xi = 0$ and $\xi = 10$ in Figure 5, even though the risk and ambiguity terms in the social discount rate decrease quadratically with η . This suggests




Figure 5: Percentage change in SE of the 2 CO₂ abatement pathway relative to BAU, as a function of the coefficient of relative risk aversion (η). The black line plots the relationship when the coefficient of relative ambiguity aversion $\xi = 0$, while the blue and red lines plot the relationship when $\xi=10$ for the recursive- and reduced-form welfare measures respectively.

that the inter-temporal substitution terms in the expression for ρ dominate the risk and ambiguity terms in our base-case simulations. For this to be so, the mean growth rate in consumption should be significantly larger than the variance of consumption growth within a given model, and the variance of mean consumption growth rates between models. This is borne out by our simulation data, for which we find inter- and intra-model variances in consumption growth to be several orders of magnitude smaller than the mean growth rate of consumption¹⁶. A further interesting feature of Figure 5 is that the difference between the recursive- and reduced-form welfare measures is small (indeed they are all but indistinguishable to the naked eye in the figure). Thus, although the assumptions these two welfare measures make about the resolution of uncertainty are poles apart, they place tight bounds on the effect of ambiguity aversion on welfare for our simulations.

One set of assumptions in our base-case simulations that deserves further scrutiny is the damage function (21), which, to recall, implies that 5°C warming will lead to damages worth the equivalent of just 6.5% of global GDP¹⁷. An increase in the global mean temperature of 5°C over the pre-industrial level is greater than the difference in temperature between the present day and the peak of the last ice age, the Last Glacial Maximum, and is thus expected to lead to biophysical changes of an unprecedented magnitude (in the context of human experience), occurring at an unprecedented rate (even in the context of geological history). Because of this, the scientific

¹⁶Neither growth rates, nor their variances, are constant in our DICE simulations, with effective annual growth rates tending to fall from approximately 2% to 1% over our 200-year simulation horizon. However the ratio of inter- and intra-model growth rate variances to the mean growth rate is smaller than 10^{-6} for every time step for the DICE damage function.

¹⁷It is important to bear in mind that total damages include both reductions in output and other reductions in welfare, such as the disutility from destruction of rare habitats, the latter estimated as an equivalent reduction in output.




Figure 6: Percentage change in SE for the 2 CO₂ abatement pathway relative to BAU, as a function of the coefficient of relative risk aversion (η), but this time with Weitzman’s damage function. The black line plots the relationship when the coefficient of relative ambiguity aversion $\xi = 0$, while the blue and red lines plot the relationship when $\xi=10$ for the recursive and reduced form welfare measures respectively.

community is essentially reduced to speculation about the consequences of around 5°C warming or more. The usual approach to calibrating the damage function is to make an estimate of the output loss accompanying 2.5°C warming or thereabouts, and then to simply extrapolate to higher temperatures based on an assumed functional form that is essentially unsupported by data of any sort. A number of scholars, including Weitzman (2010), consider predicted damages of 6.5% of GDP given 5°C warming to be remarkably low, and suggest that the damage function should be revisited. Thus we follow Weitzman (2010) in replacing (21) with a damage function which replicates its behaviour at better understood, lower temperatures, but which exhibits rapidly increasing, continuously differentiable, but threshold-like damages for higher temperatures,

$$\Omega(t) = \frac{1}{1 + (\tilde{\alpha}_1 T(t))^2 + (\tilde{\alpha}_2 T(t))^\gamma}, \quad (23)$$

where $\tilde{\alpha}_1=0.049$, $\tilde{\alpha}_2=0.16$ and $\gamma=6.75$. Like the function in (21), this choice gives damages of 1.7% of GDP for 2.5°C warming and still only 5.1% of GDP for 3.5°C warming, before increasing sharply to give damages of 9% of GDP for 4°C warming, 25% for 5°C warming, and so on. For obvious reasons Weitzman (2010, p.15) gives this particular calibration the moniker “give the devil his due”.

Figure 6 repeats the analysis of Figure 5, but this time for the more extreme damage function (23). The results are strikingly different. First, notice from the scale of the vertical axis that the Δ of the 2 CO₂ abatement policy is in general significantly larger, as we would have expected

given the often much greater damages that can be avoided. Second, and most importantly, notice that Δ is now much more sensitive to ξ . Indeed, the effect of ambiguity on welfare differences is now easy to see, with the gap between abatement and BAU policies widening substantially once ambiguity aversion is accounted for. While $\Delta_{\xi=0}$ and $\Delta_{\xi=10}$ necessarily still coincide at $\eta = 1$, the ambiguity effect is large for all other values of η we are able to include, and, in particular, increases when η rises above unity. We can again explain this result by recourse to the expression for the certainty-equivalent social discount rate in (22). With high damages, the intra- and, critically, inter-model variance in consumption growth is now larger in relation to mean growth¹⁸. Thus the third and fourth terms in (22) are larger and work against the positive first and second terms more effectively, causing ρ to decline significantly relative to its value for the DICE damage function. This gives rise to an ambiguity effect that can be very large for the higher values of η .

Finally and perhaps somewhat surprisingly, Figure 6 also shows that the ranking of the recursive and reduced form welfare measures is not independent of the value of η . When η is below approximately 2, the reduced form gives a lower estimate of Δ than the recursive form, while the opposite is true for all other values of η (except of course $\eta = 1$). In general the ranking of these measures as a function of η is highly complex, and depends on the empirical details of the consumption streams, the distribution set $\{\pi_m(S)\}$, and the value of ξ . For a heuristic explanation of the origins of the complexity of this dependence, see our discussion of a simplified model in Appendix C. Of course, the ranking of these measures is of no consequence if our primary concern is to place meaningful upper and lower bounds on the true welfare evaluation – all we need to know is that the true value lies between these two measures.

4 Discussion and conclusions

This paper aimed to provide insight into how ambiguous beliefs, and aversion to ambiguity, affect the welfare analysis of climate change abatement policy. We have argued that our knowledge of the climate system is not of sufficiently high quality to be described with unique probability distributions, and that formal frameworks that account for aversion to ambiguity are normatively legitimate, so that they provide a more accurate representation of our state of knowledge and our preferences. The paper investigates how such preferences, as represented by a smooth ambiguity model, affect optimal abatement policy in some simple static cases, and explores the conditions under which ambiguity aversion has a significant effect on the welfare evaluation of exogenously specified dynamic abatement policies.

Several lessons were learned from our static models. We showed that ambiguity aversion has a simple effect on the ambiguity-adjusted second-order probabilities on models at the level of the first-order conditions, with increases in ambiguity placing more weight on models with low expected utilities. The ambiguity-adjusted model probabilities are however endogenously determined, implying that it is not a simple matter to reason from changes in model probabilities to changes in optimal abatement. We derived conditions on the dependence of model expected utilities on abatement that are sufficient for an increase in ambiguity aversion to increase the optimal

¹⁸In the 22nd century, the ratios of inter- and intra-model spreads in consumption growth to mean consumption growth increase by 1-2 orders of magnitude for the damage function (23) relative to their values for the DICE damage function.

level of abatement. We illustrated the use and limitations of these sufficient conditions through two specific examples. In the former, we presented a model which used a simple representation of climate damages, and used our comparative statics result to define conditions under which optimal abatement increases when the decision maker’s ambiguity preferences undergo an arbitrary concave transformation. In the latter, we presented an example with a more general representation of climate damages which does not admit treatment with our general comparative statics result. Nevertheless, by specifying a particular representation of ambiguity aversion, it was possible to compute optimal abatement, and the endogenous second-order probabilities, explicitly. In this example we showed that if the expected benefits of abatement outweigh its costs, then optimal abatement is increasing in the coefficient of relative ambiguity aversion. When this benefit-cost condition is satisfied the ambiguity adjusted second-order probabilities also exhibit a simple dependence on ambiguity aversion, with their mean monotonically decreasing in the coefficient of relative ambiguity aversion.

While the static models provide an insight into how ambiguity averse preferences differ from standard risk aversion, and yield suggestive comparative statics results, their relevance to assessing realistic climate policies is limited, as they do not account for the fundamentally dynamic nature of such policies. We thus extended our analysis to the evaluation of dynamic abatement policies by using the inter-temporal version of the smooth ambiguity model. We defined two welfare measures based on extremal assumptions about the resolution of ambiguity – a recursive measure in which ambiguity is perfectly persistent, and a reduced form measure in which ambiguity resolves after a single time step – and used these to place bounds on the ambiguity averse welfare evaluation. Using the well-known DICE integrated assessment model, we computed these welfare measures for the case in which only the climate sensitivity parameter is assumed to be ambiguous. Beliefs about the climate sensitivity were represented by 20 estimated distributions from the scientific literature, and each distribution was assigned an equal second-order probability.

Using the above specification we showed that the fractional change in the stationary equivalent of abatement relative to ‘Business as usual’ is increasing in the coefficient of relative ambiguity aversion. We also investigated the interaction between risk and ambiguity preferences in determining the welfare differences between abatement and business as usual. When the coefficient of relative risk aversion (η) is 1, welfare differences are insensitive to ambiguity aversion. However, for values of η greater or less than 1 the sensitivity of welfare differences to ambiguity aversion depends significantly on the damage function. We took for our base case a ‘low’ damage function from Nordhaus (2008) and found differences in the stationary equivalent to be relatively insensitive to ambiguity aversion, due to the dominance of preferences for inter-temporal substitution. By contrast, with a ‘high’ damage function that might be considered to give more plausible output losses at high global mean temperatures (Weitzman, 2010), SE differences are very sensitive to ambiguity aversion, in particular for larger values of η . In this case, the risk and ambiguity effects dominate the inter-temporal substitution effect. What underpins both cases is the inter-model spread in mean consumption growth rates over our set of distributions for the climate sensitivity. In the base case it is small for all time. With ‘high’ damages it is substantially larger, and increases with time. Given that many commentators find values of η greater than 1 to be appropriate in the context of evaluating climate-change policies (e.g. Dasgupta, 2008; Nordhaus, 2008; Stern, 2008), we conclude that ambiguity aversion is plausibly a major driver of the economic case for abating

greenhouse gas emissions.

The results presented in this paper are subject to several limitations and qualifications. We have focussed on ambiguity over the climate sensitivity, however there are many other parameters in climate-economy models about which we may have ambiguous beliefs. The parameters of the damage function are likely candidates, as are parameters determining the cost of abatement. Therefore we should be cautious in extrapolating our findings. However, if it is supposed that there is less ambiguity about the costs of abatement than there is about the benefits (e.g. Dietz & Stern, 2008), then our qualitative finding that the value of abatement increases in ambiguity aversion may well still hold. This is because our supposition implies that an increase in abatement is more likely to reduce the spread in expected utilities across models than it is to increase it.

A further qualification surrounds our calculations in the dynamic analysis, which assumed equal second-order probabilities on each of the empirical distributions for the climate sensitivity. The smooth ambiguity model's preference representation does not provide guidance as to how these probabilities should be chosen, any more than Savage's results suggest a 'correct' prior. A possible justification for our assumption would be to appeal to a second-order principle of insufficient reason. There are however many well-known difficulties with this principle¹⁹, and we do not believe that it provides a satisfactory method for choosing second-order probabilities. Rather, these probabilities should be treated as primitive components of our beliefs, and are thus subject to rigorous elicitation. Our choice of equal probabilities reflects our inability to discern differences in the quality of the various empirical estimates of the climate sensitivity distribution – no doubt experts on these estimates would advocate alternative specifications.

At a more fundamental level, the dependence of the smooth ambiguity model's preference representation on second-order probabilities does leave us a little uneasy. Suppose that our beliefs about the veracity of the models in our set are not sufficiently well structured for us to be able to define unique second-order probabilities. Are we then committed to an infinite regress of nested preferences? While the separation between tastes and beliefs that the smooth ambiguity model achieves is a desirable property, it may not be justified in situations of deep uncertainty. The very general representation obtained in Schmeidler (1989) provides an alternative, but suffers from problems of its own²⁰. In addition, we have assumed that our model's state-space accurately describes the evolution of the climate-economy system, i.e. that there are no 'surprises'. This seems a strong assumption, given known inadequacies in our understanding of the climate system, not to mention economic aberrations. Decision theories that account for these deficiencies in our understanding have been proposed (e.g. Gilboa & Schmeidler, 1995), although it is as yet unclear whether they can be usefully applied in the climate change context.

Despite these limitations, we have shown that under certain conditions ambiguity aversion can have a significant effect on both optimal abatement, and the welfare evaluation of exogenous policies. This suggests the importance of this line of research for designing climate solutions that

¹⁹Perhaps most famously, it is not 'reparameterization invariant', a point made by Keynes (1921). This is not to say that non-informative priors that are reparameterization invariant, i.e. Jeffreys priors, would provide a satisfactory alternative.

²⁰These include the fact that the representation theorem depends on an explicit 'uncertainty aversion' axiom. One would hope that such a behavioural constraint would be an optional special case of the representation (much as risk aversion is in expected utility theory), rather than a primitive requirement. In addition, Epstein (1999) has argued that this representation is neither necessary nor sufficient to explain ambiguity averse decisions in Ellsberg-type experiments.

are true reflections of our preferences and state of knowledge.

Acknowledgments:

This research has been supported by the Munich Re programme of the Centre for Climate Change Economics and Policy; “Evaluating the Economics of Climate Risks and Opportunities”. AM gratefully acknowledges financial support from a Ciriacy-Wantrup postdoctoral fellowship.

A Proof of Proposition 2

Start with our general objective function:

$$V(a) = \sum_m p_m \phi(EU_m(a)) \quad (24)$$

where $EU_m(a)$ is the expected utility in model m when abatement is a , and m indexes the set of models $\mathcal{M}$. To prove the result, begin by considering an ordered set of functions f_λ indexed by $\lambda \in \mathbb{R}$ such that when $\lambda_2 > \lambda_1$ we have that $f_{\lambda_2} = r \circ f_{\lambda_1}$ for some concave function r , and define f_{λ_0} to be the identity function.

The first order condition for the problem (2) when ambiguity preferences are given by $f_\lambda \circ \phi$ is:

$$V'_\lambda(a) := \sum_m p_m f'_\lambda(\phi(EU_m)) \phi'(EU_m) \frac{dEU_m(a)}{da} = 0, \quad (25)$$

Note that ambiguity preferences $f_\lambda \circ \phi$ are always more ambiguity averse than ϕ when $\lambda > \lambda_0$, by definition. Let a_0 be the solution of $V'_{\lambda_0}(a_0) = 0$, i.e. the solution to the optimization problem when ambiguity preferences are given by ϕ . Now notice that $\frac{d^2 EU_m}{da^2} < 0$ implies that $\frac{d^2 V}{da^2} < 0$ ²¹. The concavity of $V(a)$ implies that, if we can find conditions under which $V'_\lambda(a_0) > 0$ for any $\lambda > \lambda_0$, then the solution of $V'_\lambda(a_\lambda) = 0$ must satisfy $a_\lambda > a_0$. Thus our strategy will be to show that the premises of the proposition imply that $V'_{\lambda_0}(a_0) = 0 \Rightarrow V'_\lambda(a_0) \geq 0$ when $\lambda > \lambda_0$.

We will make use of two lemmas to establish the result. The first is described in detail in Gollier (2001, p. 102):

Lemma 1. Let $g(x)$ be a function that crosses the x -axis singly from below (i.e. $\exists x_0$ such that, $\forall x, (x - x_0)g(x) \geq 0$). Consider a positive function $h(x, \theta)$. Then the following condition holds:

$$E_x g(x) h(x, \theta_1) = 0 \Rightarrow \forall \theta_2 > \theta_1, E_x g(x) h(x, \theta_2) \geq 0, \quad (26)$$

if and only if the function $h(x, \theta)$ is log-supermodular²², where E_x is the expectation operator over x , which is an arbitrarily distributed random variable. The same result obtains if the function $g(x)$ crosses singly from above, and the function $h(x, \theta)$ is log-submodular.

We now prove a second lemma.

²¹To see this, differentiate (2) twice to find $V''(a) = \sum_m p_m \left[\phi''(EU_m) \left(\frac{dEU_m}{da} \right)^2 + \phi'(EU_m) \frac{d^2 EU_m}{da^2} \right]$. Thus since $\phi'' < 0$ and $\phi' > 0$, $\frac{d^2 EU_m}{da^2} < 0$ implies that $V''(a) < 0$.

²²If $h(x, \theta)$ is differentiable in x , h is log-supermodular (log-submodular) if and only if $\frac{\partial h / \partial x}{h}$ is non-decreasing (non-increasing) in θ . See Gollier (2001, chap. 7) for a general definition.

Lemma 2. Define $K_a(m, \lambda) = f'_\lambda(\phi(EU_m(a)))\phi'(EU_m(a))$. $K_a(m, \lambda)$ is log-supermodular (log-submodular) when the sequence $(EU_m(a))_{m \in \mathcal{M}}$ is decreasing (increasing) in m , for any value of a .

Proof. The result is easiest to obtain in the case where m indexes a continuous set of models, and $K_a(m, \lambda)$ is differentiable in m . In this case, we have that (suppressing the a dependence of EU_m):

$$\frac{\partial K_a(m, \lambda)/\partial m}{K_a(m, \lambda)} = \left(-\frac{f''_\lambda(\phi(EU_m))}{f'_\lambda(\phi(EU_m))} \right) \phi'(EU_m) \left[-\frac{\partial EU_m}{\partial m} \right] + \frac{\phi''(EU_m)}{\phi'(EU_m)} \frac{\partial EU_m}{\partial m}. \quad (27)$$

The only term in this expression that depends on λ is the term in round brackets, which is just the coefficient of absolute ambiguity aversion for the function f_λ . So this term is increasing in the index λ , by definition²³. Thus if $\frac{\partial EU_m}{\partial m} < 0$, the expression (27) shows that $\frac{\partial K_a(m, \lambda)/\partial m}{K_a(m, \lambda)}$ is non-decreasing in λ . Similarly, when $\frac{\partial EU_m}{\partial m} > 0$, (27) shows that $\frac{\partial K_a(m, \lambda)/\partial m}{K_a(m, \lambda)}$ is non-increasing in λ . Similar, if slightly more involved, computations establish the result in the case of a discrete set of models. $\square$

We now combine these two lemmas. Consider the case in which $(\frac{dEU_m}{da})_{m \in \mathcal{M}}$ is increasing in m , and $(EU_m)_{m \in \mathcal{M}}$ is decreasing in m . Define

$$g(m; a_0) := \frac{dEU_m}{da} \Big|_{a=a_0}. \quad (28)$$

Evaluating the first order condition (25) at $\lambda = \lambda_0$, we have that

$$\sum_m p_m \phi'(EU_m(a_0)) \frac{dEU_m}{da} \Big|_{a=a_0} = 0. \quad (29)$$

Since $(\frac{dEU_m}{da})_{m \in \mathcal{M}}$ is increasing in m , (29) makes it clear that $g(m; a_0)$ crosses the horizontal axis singly from below. This is so since $\phi' > 0$ means that in order for (29) to be satisfied, we require some of the terms in the sequence $(\frac{dEU_m}{da} \Big|_{a=a_0})_{m \in \mathcal{M}}$ to be negative, and some to be positive. Since the sequence is increasing, $g(m; a_0)$ must exhibit the single-crossing property.

Now since $\{EU_m\}$ is decreasing in m by assumption, by Lemma 2 we have that $K_{a_0}(m, \lambda)$ is a log-supermodular function. By the definition of $g(m; a_0)$ and λ_0 ,

$$\sum_m p_m g(m; a_0) K_{a_0}(m, \lambda_0) = V'_{\lambda_0}(a_0) = 0. \quad (30)$$

Now by Lemma 1, the fact that $g(m; a_0)$ exhibits the single crossing property, and the fact that $K_{a_0}(m, \lambda)$ is log-supermodular, we know that for any $\lambda > \lambda_0$,

$$\sum_m p_m g(m; a_0) K_{a_0}(m, \lambda) = V'_\lambda(a_0) \geq 0, \quad (31)$$

establishing the result.

²³Note that for concave functions f_1, f_2 , the coefficient of absolute risk aversion (CARA) for f_2 is uniformly larger than the CARA for f_1 iff $f_2 = r \circ f_1$ for some concave r .

To complete the proof, notice that when $(EU_m)_{m \in \mathcal{M}}$ is increasing and $\left(\frac{dEU_m}{da}\big|_{a=a_0}\right)_{m \in \mathcal{M}}$ decreasing in m , we can apply the same reasoning and use Lemma 2 in the case of a function g that crosses singly from above, and a function h that is log-submodular.

Finally, this method of proof is easily extended to the case where the sequences $(EU_m)_{m \in \mathcal{M}}$ and $\left(\frac{dEU_m}{da}\big|_{a=a_0}\right)_{m \in \mathcal{M}}$ are comonotonic, in which case ambiguity aversion has the opposite affect on optimal abatement. To do this, one simply defines $g(m; a_0) := -\frac{dEU_m}{da}\big|_{a=a_0}$ in the step (28). Under the assumption that both sequences are decreasing in m for all a , all the following steps of the proof go through unchanged, and we are left with the conclusion that $-V'_{\lambda_0}(a_0) = 0 \Rightarrow -V'_\lambda(a_0) \geq 0$ when $\lambda > \lambda_0$. Thus by the concavity of $V(a)$, an increase in λ *decreases* the optimal value of a in this case.

B Calculations for Example 2

$$V(a) = \int p(\theta) \phi \left(\int U(c_0 + (1 - b_1 a)x - b_2 a) \pi_\theta(x) dx \right) d\theta.$$

where we assume that

$$\begin{aligned} p(\theta) &\sim \mathcal{N}(\mu, \sigma_0) \\ \pi_\theta(x) &\sim \mathcal{N}(\theta, \sigma) \\ U(c) &= -\frac{1}{A} \exp(-Ac) \\ \phi(U) &= -(-U)^{1+\xi} / (1 + \xi). \end{aligned}$$

Using the fact that

$$\int e^{\lambda x} \pi(x) dx = e^{\lambda(\theta + \frac{1}{2}\lambda\sigma^2)}$$

for any constant λ , and any $\pi(x) \sim \mathcal{N}(\theta, \sigma)$, one can compute all the integrals in the expression for $V(a)$ to find that

$$V(a) \propto -\exp \left(-A(1 + \xi) \left[c_0 - b_2 a + (1 - b_1 a)\mu - \frac{1}{2} A(1 - b_1 a)^2 (\sigma^2 + (1 + \xi)\sigma_0^2) \right] \right).$$

$V(a)$ is maximized when the term in square brackets above is maximized. The solution is

$$a^* = \frac{1}{b_1} + \frac{b_1 \mu + b_2}{A b_1 (\sigma^2 + (1 + \xi)\sigma_0^2)}.$$

Now the ambiguity adjusted distribution of the means θ is:

$$\hat{p}(\theta) = \frac{p(\theta) \phi' \left(\int U(c_0 + (1 - b_1 a^*)x - b_2 a^*) \pi_\theta(x) dx \right)}{\int p(\theta) \phi' \left(\int U(c_0 + (1 - b_1 a^*)x - b_2 a^*) \pi_\theta(x) dx \right) d\theta}$$

Consider the numerator. $\phi'(U) = (-U)^\xi$, so this is just

$$A^{-\xi} p(\theta) \exp \left(-A\xi \left[c_0 - b_2 a^* + (1 - b_1 a^*)\theta - \frac{1}{2} A\sigma^2 (1 - b_1 a^*)^2 \right] \right).$$

All the terms that do not depend on θ are going to get divided out, so the important part of the numerator is

$$p(\theta) \exp(-A\xi(1 - b_1 a^*)\theta).$$

Because $p(\theta)$ is a normal distribution, we can multiply the exponential into the normal PDF. Once again, we can neglect everything that does not depend on θ , as it will get divided out:

$$\begin{aligned} p(\theta) \exp(-A\xi(1 - b_1 a^*)\theta) &\propto \exp\left(-\frac{(\theta - \mu)^2}{2\sigma_0^2}\right) \exp(-A\xi(1 - b_1 a^*)\theta) \\ &\propto \exp\left(-\frac{1}{2\sigma_0^2}[\theta^2 - 2\theta(\mu - A\sigma_0^2\xi(1 - b_1 a^*))]\right) \\ &\propto \exp\left(-\frac{[\theta - (\mu - \sigma_0^2\xi A(1 - b_1 a^*))]^2}{2\sigma_0^2}\right) \end{aligned}$$

In the last step we have just completed the square. This expression looks like another normal PDF, only with an adjusted mean. Since we know that $\hat{p}(\theta)$ will be properly normalized, we can conclude that

$$\hat{p}(\theta) \sim \mathcal{N}(\mu - \sigma_0^2\xi A(1 - b_1 a^*), \sigma_0^2).$$

C Ranking recursive and reduced form welfare measures

In order to gain an intuition for why the ranking of recursive and reduced form welfare measures has a complex dependence on η , consider a 3-period model, and assume that beliefs are described by only two different probability models. Let $U(0)$ be the known utility at time 0, $EU_i(t)$ be the expected utility in model $i \in \{1, 2\}$ at time $t \in \{1, 2\}$, and p be the second-order probability on distribution 1. The two welfare measures are given by:

$$\begin{aligned} V_{reduced} &= U(0) + \beta\phi^{-1}(p\phi(EU_1(1) + \beta EU_1(2)) + (1 - p)\phi(EU_2(1) + \beta EU_2(2))) \\ V_{recursive} &= U(0) + \beta\phi^{-1}(p\phi(EU_1(1) + \beta X) + (1 - p)\phi(EU_2(1) + \beta X)) \end{aligned}$$

where

$$X := \phi^{-1}(p\phi(EU_1(2)) + (1 - p)\phi(EU_2(2))).$$

From these expressions, $V_{recursive} > V_{reduced}$ iff

$$\mathbf{E}_i\phi(EU_i(1) + \beta X) > \mathbf{E}_i\phi(EU_i(1) + \beta EU_i(2)),$$

i.e., we are asking whether adding a certainty equivalent X of the risk $EU_i(2)$ to the risk $EU_i(1)$ is preferred to adding the risk $EU_i(2)$ itself. There is no general answer to this question – it depends on the values of the $EU_i(t)$, and the specific functional form of ϕ . Moreover, since the $EU_i(t)$ are all implicitly complicated functions of η , and are likely non-monotonic, and discontinuous at $\eta = 1$, we can expect to get ranking reversals as η changes and affects the values of the $EU_i(t)$ non-uniformly.

References

- M. Allen, et al. (2006). ‘Observational constraints on climate sensitivity’. In H. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, & G. Yohe (eds.), *Avoiding dangerous climate change*, p. 406. Cambridge University Press, Cambridge, UK.
- F. J. Anscombe & R. J. Aumann (1963). ‘A Definition of Subjective Probability’. *The Annals of Mathematical Statistics* **34**(1):199–205.
- K. Arrow & L. Hurwicz (1977). ‘An optimality criterion for decision-making under ignorance.’. In *Studies in resource allocation processes*, p. 482. Cambridge University Press, Cambridge, UK.
- G. W. Bassett, et al. (2004). ‘Pessimistic Portfolio Allocation and Choquet Expected Utility’. *Journal of financial econometrics* **2**(4):477–492.
- K. Binmore (2009). *Rational Decisions*. Princeton University Press.
- P. Bossaerts, et al. (2010). ‘Ambiguity in Asset Markets: Theory and Experiment’. *Rev. Financ. Stud.* p. 106.
- C. Camerer (1999). ‘Ambiguity-aversion and non-additive probability: Experimental evidence, models and applications’. In L. Luini (ed.), *Uncertain decisions: bridging theory and experiments*, pp. 53–80. Kluwer academic publishers.
- P. Dasgupta (2008). ‘Discounting climate change’. *Journal of Risk and Uncertainty* **37**(2):141–169.
- S. Dietz & N. Stern (2008). ‘Why Economic Analysis Supports Strong Action on Climate Change: A Response to the Stern Review’s Critics’. *Review of Environmental Economics and Policy* **2**(1):94–113.
- J. Dow & S. R. da Costa Werlang (1992). ‘Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio’. *Econometrica* **60**(1):197–204.
- D. Ellsberg (1961). ‘Risk, Ambiguity, and the Savage Axioms’. *The Quarterly Journal of Economics* **75**(4):643–669.
- L. G. Epstein (1999). ‘A Definition of Uncertainty Aversion’. *Review of Economic Studies* **66**(3):579–608.
- D. J. Frame, et al. (2007). ‘Probabilistic climate forecasts and inductive problems’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):1971–1992.
- P. Ghirardato, et al. (2004). ‘Differentiating ambiguity and ambiguity attitude’. *Journal of Economic Theory* **118**(2):133–173.
- I. Gilboa (2009). *Theory of Decision under Uncertainty*. Cambridge University Press, 1 edn.
- I. Gilboa, et al. (2009). ‘Is It Always Rational to Satisfy Savage’s Axioms?’. *Economics and Philosophy* **25**(3):285–296.

- I. Gilboa, et al. (2008). ‘Probability and Uncertainty in Economic Modeling’. *Journal of Economic Perspectives* **22**(3):173–188.
- I. Gilboa & D. Schmeidler (1989). ‘Maxmin expected utility with non-unique prior’. *Journal of Mathematical Economics* **18**(2):141–153.
- I. Gilboa & D. Schmeidler (1995). ‘Case-Based Decision Theory’. *The Quarterly Journal of Economics* **110**(3):605–39.
- C. Gollier (2001). *The economics of risk and time*. MIT Press, Cambridge, Mass. ; London.
- C. Gollier (2009). ‘Portfolio choices and asset prices: The comparative statics of ambiguity aversion, <http://idei.fr/display.php?a=4812>’. *IDEI Working Paper* (357).
- C. Gollier & J. Gierlinger (2008). ‘Socially efficient discounting under ambiguity aversion, <http://idei.fr/display.php?a=9848>’. *IDEI Working Paper* (561).
- L. P. Hansen & T. J. Sargent (2007). *Robustness*. Princeton University Press.
- G. Heal (2009). ‘Climate Economics: A Meta-Review and Some Suggestions for Future Research’. *Rev Environ Econ Policy* **3**(1):4–21.
- C. Henry & M. Henry (2002). ‘Formalization and application of the precautionary principle, <http://www.columbia.edu/cu/economics/discapwr/DP0102-22.pdf>’. *Columbia University Department of Economics Discussion Paper Series*.
- C. Hope (2006). ‘The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC’s five reasons for concern’. *Integrated Assessment* **6**(1).
- N. Ju & J. Miao (2009). ‘Ambiguity, learning, and asset returns’. *Working paper* (<http://people.bu.edu/miao/jm17.pdf>).
- K. Keller, et al. (2004). ‘Uncertain climate thresholds and optimal economic growth’. *Journal of Environmental Economics and Management* **48**(1):723–741.
- D. L. Kelly & C. D. Kolstad (1999). ‘Bayesian learning, growth, and pollution’. *Journal of Economic Dynamics and Control* **23**(4):491–518.
- J. M. Keynes (1921). *A treatise on probability*. Macmillan and Co.
- P. Klibanoff, et al. (2005). ‘A Smooth Model of Decision Making under Ambiguity’. *Econometrica* **73**(6):1849–1892.
- P. Klibanoff, et al. (2009). ‘Recursive smooth ambiguity preferences’. *Journal of Economic Theory* **144**(3):930–976.
- F. Knight (1921). *Risk, uncertainty, and profit*. Houghton Mifflin, New York.
- A. Lange & N. Treich (2008). ‘Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions’. *Climatic Change* **89**(1):7–21.

- A. Manne & R. Richels (1992). *Buying Greenhouse Insurance: The Economic Costs of CO₂ Emission Limits*. The MIT Press.
- M. Meinshausen, et al. (2009). ‘Greenhouse-gas emission targets for limiting global warming to 2C’. *Nature* **458**(7242):1158–1162.
- W. D. Nordhaus (2008). *A question of balance*. Yale University Press.
- W. A. Pizer (1999). ‘The optimal choice of climate change policy in the presence of uncertainty’. *Resource and Energy Economics* **21**(3-4):255–287.
- L. J. Savage (1954). *The Foundations of Statistics*. Wiley and Sons.
- D. Schmeidler (1989). ‘Subjective Probability and Expected Utility without Additivity’. *Econometrica* **57**(3):571–587.
- P. Slovic & A. Tversky (1974). ‘Who accepts Savage’s axiom?’. *Behavioral Science* **19**(6):368–373.
- L. A. Smith (2002). ‘What might we learn from climate forecasts?’. *Proceedings of the National Academy of Sciences of the United States of America* **99**(Suppl 1):2487–2492.
- L. A. Smith (2007). *Chaos : a very short introduction*, vol. 159. Oxford University Press, Oxford.
- D. Stainforth, et al. (2007). ‘Confidence, uncertainty and decision-support relevance in climate predictions’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):2145–2161.
- N. Stern (2008). ‘The Economics of Climate Change’. *American Economic Review* **98**(2):1–37.
- R. Tol (1997). ‘On the optimal control of carbon dioxide emissions: an application of FUND’. *Environmental Modeling and Assessment* **2**(3):151–163.
- C. P. Traeger (2009). ‘Recent Developments in the Intertemporal Modeling of Uncertainty’. *Annual Review of Resource Economics* **1**(1):261–286.
- J. von Neumann & O. Morgenstern (1944). *Theory of Games and Economic Behaviour*. Princeton University Press.
- M. L. Weitzman (1976). ‘On the Welfare Significance of National Product in a Dynamic Economy’. *The Quarterly Journal of Economics* **90**(1):156–162.
- M. L. Weitzman (2009). ‘On Modeling and Interpreting the Economics of Catastrophic Climate Change’. *Review of Economics and Statistics* **91**(1):1–19.
- M. L. Weitzman (2010). ‘GHG targets as insurance against catastrophic climate damages’. *Working paper* .