

Pollution control: when, and how, to be precautionous

Stergios Athanassoglou*, Anastasios Xepapadeas†

January 21, 2011

Abstract

The precautionary principle (PP) applied to environmental policy stipulates that, in the presence of physical uncertainty, society must take take robust preventive action to guard against worst-case outcomes. It follows that the higher the degree of uncertainty, the more aggressive this preventive action should be. This normative maxim is explored in the case of a stylized dynamic model of pollution control under Knightian uncertainty. At time 0 a decisionmaker makes a one-time investment in damage-control technology and subsequently decides on a desirable dynamic emissions policy. Adopting the robust control framework of Hansen and Sargent [10], we investigate optimal damage-control and mitigation policies. We show that optimal investment in damage control is always increasing in the degree of uncertainty, thus confirming the conventional PP wisdom. Optimal mitigation decisions, however, need not always comport with the PP and we provide analytical conditions that sway the relationship one way or the other. This result is interesting when contrasted to a model with fixed damage-control technology, in which it can be easily shown that a PP vis-a-vis mitigation unambiguously holds. We conduct a set of numerical experiments to determine the sensitivity of our results to specific functional forms of damage-control cost. We find that when the cost of damage-control technology is low enough, damage-control investment and mitigation may act as substitutes and a PP with respect to the latter can be unambiguously irrational.

Keywords: risk, ambiguity, robust control, precautionary principle, pollution control

JEL classifications: C61, D80, D81

*Fondazione Eni Enrico Mattei and Euro-Mediterranean Center for Climate Change, athanassoglou@gmail.com.

†Athens University of Economics and Business and Beijer Fellow, xepapad@aueb.gr. We thank Valentina Bosetti and seminar participants at the Fondazione Eni Enrico Mattei for useful comments.

1 Introduction

A common thread running through much of environmental economics is a reliance on expected utility as a means of performing cost-benefit analysis and, more broadly, as a normative criterion. Expected utility theory has solid theoretical underpinnings, going back to the groundbreaking work of von Neumann and Morgenstern [17] and Savage [21], and leads to tractable optimization problems. However, in the case of environmental economics, its attractive qualities often come at a steep price, primarily due to two basic factors: (a) the high structural uncertainty over the physics of environmental phenomena which makes the assignment of precise probabilistic model structure untenable, and (b) the high sensitivity of model outputs to seemingly ad hoc modeling assumptions (for instance, the functional form of the chosen damage function and the value of the social discount rate) on which often little consensus exists. As a result, separate models may arrive at dramatically different policy recommendations, generating heated debate and much confusion over the magnitude and timing of desirable policy.¹

Weitzman [26] has forcefully made the above point in the context of climate-change economics. Echoing the previous discussion, his main thesis is that the deep structural uncertainty surrounding the science of climate change creates substantial problems for the application of classical cost-benefit analysis. His assessment rests on the incontestable premise that climate scientists are far from reaching consensus on the physical effects of unprecedented concentrations of GHGs in the atmosphere. In particular, state-of-the-art climate science cannot rule out the occurrence of extreme catastrophic events with nontrivial probability. In addition, Weitzman [27], as well as others (Stern and Persson [23]) argue that the results of established economic models (such as Nordhaus' RICE and DICE models [18]) are too sensitive to the precise way in which damages are mathematically defined, which constitutes yet another modeling convention that cannot, at present, be scientifically validated to an adequate degree of confidence. The main implication of this line of research is that, in the presence of such extreme uncertainty, it is far from clear how one should model the economic implications of climate change policy. Indeed, Weitzman himself admits to no grand theoretical panacea for this impasse and merely concedes that his analysis calls, in a qualitative sense, for extremely cautious policy-making. While the academic community has been reluctant to embrace the totality of his analysis (Nordhaus [19]; Karp [13]), Weitzman's work goes a long way in illuminating the potential pitfalls of discounting the likelihood of extreme climatic events and cavalierly insisting on a particular modeling framework.

Weitzman's criticism notwithstanding, environmental policy needs to be predicated on a coherent, if simplified, account of its economic implications. We focus our attention on the problem of deep and structural uncertainty that is defined by an inability to posit precise probabilistic structure to physical processes. This is close to the concept of uncertainty as introduced by Knight [15] to represent a situation where there is ignorance, or not enough information, to assign probabilities to

¹William Nordhaus' DICE model [18] and the Stern Report [22] are prominent examples of this deep divergence.

events. Knight argued that this deeper kind uncertainty is quite common in economic decision making, and thus deserving of systematic study. Knightian uncertainty is contrasted to risk (measurable or probabilistic uncertainty) where probabilities can be assigned to events and are summarized by a subjective probability measure or a *single* Bayesian prior.

Inspired by the the work of Knight [15] and consequently Ellsberg [4], economic theorists have questioned the classical expected utility framework and attempted to formally model preferences in environments in which probabilistic beliefs are not of sufficiently high quality to generate prior distributions. Klibanoff et al. [14] developed an axiomatic framework, the so-called “smooth ambiguity” model, in which different degrees of aversion for uncertainty are explicitly parameterized in agents’ preferences. In their model an act f is preferred to an act g if and only if $\mathbf{E}_p\phi(\mathbf{E}_\pi u \circ f) > \mathbf{E}_p\phi(\mathbf{E}_\pi u \circ g)$, where u is a von Neumann Morgenstern utility function, ϕ an increasing function, and p is a subjective second order probability over a set Π of probability measures π that the decisionmaker is willing to consider (\mathbf{E} denotes the expectation operator). When ϕ is concave the decisionmaker is said to be *ambiguity averse*. This theoretical framework has been recently applied to environmental economics, particularly in issues related to climate change (Traeger [24]; Gollier and Gierlinger [8]; Millner et al. [16]). However, despite its notable role in the recent literature, the smooth ambiguity model seems to have more of a positive instead of a normative focus, and questions about how to calibrate agents’ ambiguity aversion seem difficult to address. An additional, potential, shortcoming of the general approach is that it relies on knowledge of second-order probabilities (the distribution p) when, in some instances, such knowledge may not be possible or justified. Moreover, in the context of applied work, the smooth ambiguity model seems to pose nontrivial tractability challenges [16] so that only the utility of very simple, exogenously given, policies can be computed.

Our Focus: Robust Control. In a seminal contribution, Gilboa and Schmeidler [7] developed the axiomatic foundations of *max-min* expected utility, a substitute of classical expected utility for economic environments featuring unknown risk. They argued that when the underlying uncertainty of an economic system is not well understood, it is sensible, and axiomatically coherent, to optimize over the worst-case outcome (i.e. the worst-case *prior*) that may conceivably occur. Doing so guards against possible devastating losses in any possible state of the world and thus adds an element of robustness to the decision-making process.² Motivated by the very real possibility of model misspecification in macroeconomics, Hansen and Sargent [10] and Hansen et al. [11] extend Gilboa and Schmeidler’s insight to continuous-time dynamic optimization problems, introducing the concept of robust control to economic environments. They show how standard dynamic programming techniques can be modified to yield robust solutions to problems in which the underlying

²A simple way to think of this idea is the following: If a company is building an airplane, its primary aim is to make sure the plane never malfunctions, no matter what (realistic) weather conditions it may face. This concern for robustness trumps the possible suboptimality of the plane’s functions (speed, efficiency, to name a few) for what are considered to be “average” conditions.

stochastic nature of the model is not perfectly known. In their work, the degree to which a model is misspecified is a model input, so that decision makers can test the sensitivity of a proposed solution with respect to the model’s presumed uncertainty. Eschewing complex formal characterizations of their results similar to Klibanoff et al. [14], the focus of their robustness project is as much practical as it is theoretical, if not more.

Finally, we should note that Chen and Epstein [2] and Epstein and Schneider [5] developed a parallel approach to Hansen and Sargent’s robust control, which they refer to as the Recursive Multiple Priors (RMP) model. Similarly inspired by Gilboa and Schmeidler, this framework differs in subtle ways to robust control, primarily with regard to the set of restricted priors, and their evolution over time. An additional substantive difference is that while robust control provides a method of computing a max-min solution, recursive multiple priors seem to function primarily as a means of comparing the minimum utility of given policies.³ A recent application of RMP in environmental economics can be found in Asano [1], who studies the optimal timing of environmental policy under ambiguity. Asano’s analysis deals with stock pollutants and verifies the existence of a precautionary principle (PP), in which an increase in ambiguity implies an earlier adoption of policy.

Our contribution. In recent years the Hansen-Sargent framework has slowly begun to make its way into environmental economics. Gonzales [9] applied robust control to the regulation of a stock pollutant under multiplicative uncertainty (first introduced by Hoel and Karp [12]). Roseta-Palma and Xepapadeas [20] studied water management under ambiguity, while Vardas and Xepapadeas [25] did the same in the context of biodiversity management. Both these contributions focused on determining the “cost of precaution,” that is, the decrease in utility that model misspecification leads to.

The present work can be viewed as a continuation of this nascent literature in the context of pollution control. Our paper expands the standard linear-quadratic model of Dockner and Van Long [3] to allow for (a) model misspecification and (b) the possibility of investment in *damage-control* technology that alleviates the effects of stock pollutant accumulation. In the context of climate change, examples of this kind of damage-control investment can be found in (a) the construction of large-scale civil engineering projects (b) substantial R & D in geoengineering and (c) the construction of new urban environments to accommodate potential forced migration. We assume the presence of a benevolent government (or, conversely, a group of cooperating countries in a global pollution control problem) which makes a one-time investment in damage-control technology at time 0, and subsequently decides on a desirable dynamic emissions policy. Adopting the Hansen-Sargent robust control framework, we introduce Knightian uncertainty into the model and study the effect of model misspecification on optimal mitigation and damage-control decisions. Unlike the aforementioned contributions, we are able to completely characterize the worst-case pollution

³For more details the reader is referred to section 5 in Epstein and Schneider [5] and section 9 in Hansen et al. [11].

accumulation process and attach a physically meaningful parameter (entropy bound) on the degree of model misspecification.

Our primary focus is normative. Ex-ante, one would expect a certain kind of precautionary principle (PP) to hold whereby, the *greater* the degree of uncertainty, the *more* the government would choose to *both* mitigate and invest in damage control. Indeed, since higher uncertainty translates to the possibility of higher damages from pollutant accumulation, such a finding would be, more or less, congruent to our intuition.

But intuitive reasoning can often be (partially) wrong. We show that optimal investment in damage control technology is always increasing in the degree of uncertainty, thus confirming the conventional PP wisdom. Optimal mitigation decisions, however, need not always comport with the PP and we provide analytical conditions that sway the relationship one way or the other. This result is interesting when contrasted to a model without the possibility for damage control, in which it can be easily shown that the PP unambiguously holds. Focusing on a tractable family of damage-control cost functions, we conduct a set of numerical experiments to determine the sensitivity of our results to specific functional forms. We find that when the cost of damage control is low enough, damage-control investment and mitigation may act as substitutes and a PP with respect to the latter can be unambiguously irrational.

Paper outline. The structure of the paper is as follows. Section 2 introduces the robust control model, while Section 3 analyzes its solution for the case in which damage-control technology is fixed. Section 4 introduces the possibility of damage-control investment and studies the applicability of a PP with respect to both mitigation and damage control. Section 5 illustrates our theoretical results with a numerical exercise that considers a set of different damage-control cost functions. Section 6 provides concluding remarks.

2 Robust Pollution Control

2.1 Introducing model misspecification and damage control technology

We employ the standard linear quadratic model of international pollution control analyzed by Dockner and van Long [3]. Output is a function of emissions $Q = F(E)$, where $F(\cdot)$ is strictly concave with $F(0) = 0$. Emissions contribute to the stock of a global pollutant $P(t)$. The evolution of the pollution stock is described by the usual linear differential equation,

$$\dot{P}(t) = E - mP(t), P(0) = P_0, \quad (1)$$

where $m > 0$ reflects the environment's self cleaning capacity. Utility, assuming constant population normalized to one, is $u(F(E)) - D(P)$ where $D(P)$ is a damage function and

$$u(F(E)) = -\frac{1}{2}E^2 + AE. \quad (2)$$

We modify the standard quadratic damage function $D(P) = sP^2$, $s > 0$ by allowing the possibility of investment in damage control. That is, at time 0, the government chooses a level of damage-control technology $z \in [0, 1]$ that alters the damage function in the following way

$$D(P, z) = z \cdot sP^2. \quad (3)$$

Thus, a lower level of z implies a higher investment in damage-control technology. The cost of making an investment z is modeled by a strictly decreasing and convex function $\phi(z) : [0, 1] \mapsto \mathfrak{R}^+$ that satisfies

$$\phi(1) = 0, \quad \lim_{z \rightarrow 0} \phi(z) = \infty, \quad \lim_{z \rightarrow 0} \phi'(z) = -\infty.$$

Possible candidates for $\phi(z)$ include $1/z^k - 1$, for $k > 0$.

Risk is introduced to the standard model so that the stock of the pollutant accumulates according to the diffusion process

$$dP(t) = (E - mP(t)) dt + \sigma dB(t), \quad (4)$$

where $\{B(t) : t \geq 0\}$ is a Brownian motion on an underlying probability space (Ω, \mathcal{F}, G) . Thus, in a world without uncertainty, the government's objective is to maximize welfare or

$$\begin{aligned} \max_E \quad & \mathbf{E} \int_0^\infty e^{-\rho t} \left[AE - \frac{E^2}{2} - szP^2 \right] dt \\ \text{subject to:} \quad & (4), \quad P(0) = P_0. \end{aligned} \quad (5)$$

Optimization problem (5) is referred to as the *benchmark* model.

If there were no fear of model misspecification solving the benchmark problem (5) would be sufficient. As this is not the case, following Hansen and Sargent [10], model misspecification can be reflected by a family of stochastic perturbations to the Brownian motion so that the probabilistic structure implied by SDE (4) is distorted and the probability measure G is replaced by another Q . The perturbed model is obtained by performing a change of measure and replacing $B(t)$ in Eq. (4) by

$$\hat{B}(t) + \int_0^t v(s) ds, \quad (6)$$

where $\{\hat{B}(t) : t \geq 0\}$ is a Brownian motion and $\{v(t) : t \geq 0\}$ is a measurable drift distortion such that $v(t) = v(P(s) : s \leq t)$. Thus, changes to the distribution of $B(t)$ are parameterized as drift distortions to a fixed Brownian motion $\{\hat{B}(t) : t \geq 0\}$. The distortions will be zero when $v \equiv 0$ and the two measures G and Q coincide. Thus pollution dynamics under model misspecification can be written as:

$$dP(t) = (E - mP(t) + \sigma v(t)) dt + \sigma dB(t). \quad (7)$$

As discussed in Hansen and Sargent [10], the discrepancy between the two measures G and Q is measured through their relative entropy

$$R(Q) = \int_0^\infty e^{-\rho t} \frac{1}{2} \mathbf{E}_Q [v(t)^2] dt, \quad (8)$$

where \mathbf{E} denotes the expectation operator. To express the idea that even when the model is misspecified the benchmark model remains a “good” approximation, the misspecification error is constrained so that we only consider distorted probability measures Q such that

$$R(Q) = \int_0^\infty e^{-\rho t} \frac{1}{2} \mathbf{E}_Q[v(t)^2] dt \leq \eta < \infty, \quad (9)$$

where $e^{-\rho t}$ is the appropriate discount factor. By modifying the value of η in (9) the decisionmaker can control the degree of model misspecification he is willing to consider. In particular, if the decisionmaker can use physical principles and time series analysis to formulate bounds on the relative entropy of plausible probabilistic deviations from his benchmark model, these bounds can be used to calibrate the parameter η .

2.2 Robust control

Under model misspecification benchmark pollution dynamics (4) are replaced by (7). Two robust control problems can be associated with the solution to the misspecified problem: (a) a *constraint* robust control problem which explicitly models a bound on relative entropy, and (b) a *multiplier* robust control problem which incorporates a Lagrange multiplier to a relative entropy constraint.

Formally, the multiplier robust control problem is defined as

$$\begin{aligned} V(P_0; \theta, z) = \max_E \min_v & \quad \mathbf{E} \int_0^\infty e^{-\rho t} \left[AE - \frac{E^2}{2} - szP^2 + \frac{\theta v^2}{2} \right] dt \\ \text{subject to:} & \quad (7), \quad P(0) = P_0, \end{aligned} \quad (10)$$

while the constraint robust control problem is given by

$$\begin{aligned} V(P_0; \eta, z) = \max_E \min_v & \quad \mathbf{E} \int_0^\infty e^{-\rho t} \left[AE - \frac{E^2}{2} - szP^2 \right] dt \\ \text{subject to:} & \quad (7), \quad (9), \quad P(0) = P_0. \end{aligned} \quad (11)$$

In both extremization problems, the distorting process v_t is such that allowable measures Q have finite entropy. In the constraint problem (11), the parameter η is the maximum expected misspecification error that the decision-maker is willing to consider. In the multiplier problem (10), the parameter θ can be interpreted as a Lagrange multiplier associated with entropy constraint $R(Q) \leq \eta$. Our choice of θ lies in an interval $(\underline{\theta}, +\infty]$, where the lower bound $\underline{\theta}$ is a breakdown point beyond which it is fruitless to seek more robustness. This is because the minimizing agent is sufficiently unconstrained so that he can push the criterion function to $-\infty$ despite the best response of the maximizing agent. Thus when $\theta < \underline{\theta}$, robust control rules cannot be attained. On the other hand when $\theta \rightarrow \infty$ or, equivalently $\eta = 0$, there are no concerns about model misspecification and the decision-maker may safely consider just the benchmark model.

The relationship between the two robust control problems is subtle. For instance, a particular θ can be associated with no, or even multiple, η 's, while a particular η can map to multiple θ 's.⁴

⁴For details the reader is referred to Sections 5 and 7 in Hansen et al. [11].

In what follows, we will focus on the multiplier problem (10). We do so because it is the more analytically tractable problem of the problem of the two (Fleming and Souganidis [6]). However, it is worth noting that, in contrast to previous contributions, our subsequent analysis is capable of providing a connecting thread to the more intuitive, and physically meaningful, constraint formulation. This is because we are able to explicitly characterize the worst-case perturbed probability measure Q^* of a given multiplier problem, to which we then apply Proposition 2 in Hansen and Sargent [10], which establishes the following:

Proposition 1 (Prop. 2, Hansen and Sargent [10]) *Suppose V is strictly decreasing in η , $\theta^* \in (\underline{\theta}, +\infty]$, and that there exists a solution E^* and v^* (corresponding to measure Q^*) to the multiplier problem (10). Then, that E^* also solves the constraint problem (11) for $\eta = \eta^* = R(Q^*)$.*

3 Robust pollution control with fixed damage control technology

3.1 Problem solution

We initially focus on solving the multiplier problem (10) for a given level of damage control technology $z \in [0, 1]$. The Bellman-Isaacs condition (see Fleming and Souganidis [6]) is given by the following equation:

$$\rho V = \max_E \min_v \left\{ AE - \frac{E^2}{2} - szP^2 + \frac{\theta v^2}{2} + V_P(E - mP + \sigma v) + \frac{\sigma^2}{2} V_{PP} \right\} \quad (12)$$

Minimizing first with respect to v , we obtain

$$v^* = -\frac{\sigma V_P}{\theta},$$

so that Eq. (12) becomes

$$\rho V = \max_E \left\{ AE - \frac{E^2}{2} - szP^2 + V_P(E - mP) + \frac{\sigma^2}{2} V_{PP} - \frac{\sigma^2}{2\theta} (V_P)^2 \right\}. \quad (13)$$

Maximizing with respect to E , we have

$$E^* = A + V_P$$

so that the differential equation we need to solve is the following

$$\rho V = A(A + V_P) - \frac{(A + V_P)^2}{2} - szP^2 + V_P(A + V_P - mP) + \frac{\sigma^2}{2} V_{PP} - \frac{\sigma^2}{2\theta} (V_P)^2. \quad (14)$$

One can show that the value function satisfying (14) admits the following simple quadratic form⁵

$$V(P; \theta, z) = \alpha_1(\theta, z)P^2 + \alpha_2(\theta, z)P + \alpha_3(\theta, z), \quad (15)$$

where

$$\alpha_1(\theta, z) = \frac{2m + \rho - \sqrt{(2m + \rho)^2 + 8sz(1 - \frac{\sigma^2}{\theta})}}{4(1 - \frac{\sigma^2}{\theta})} \leq 0 \quad (16)$$

$$\alpha_2(\theta, z) = \frac{2A\alpha_1(\theta, z)}{2\alpha_1(\theta, z)(\frac{\sigma^2}{\theta} - 1) + \rho + m} \leq 0 \quad (17)$$

$$\begin{aligned} \alpha_3(\theta, z) &= \frac{1}{\rho} \left[A^2 + \alpha_2(\theta, z)A - \frac{(A + \alpha_2(\theta, z))^2}{2} + \alpha_2(A + \alpha_2(\theta, z)) + \sigma^2\alpha_1(\theta, z) - \frac{\sigma^2}{2\theta}\alpha_2(\theta, z)^2 \right] \\ &= \frac{1}{\rho} \left[\frac{A^2}{2} + \sigma^2\alpha_1(\theta, z) + \alpha_2(\theta, z)A + \frac{\alpha_2(\theta, z)^2}{2} \left(1 - \frac{\sigma^2}{\theta}\right) \right] \end{aligned} \quad (18)$$

Max-min optimal emissions E^* satisfy

$$E^*(P, \theta, z) = A + V_P = A + \alpha_2(\theta, z) + 2\alpha_1(\theta, z)P, \quad (19)$$

and the worst-case misspecification v^* is given by

$$v^*(P, \theta, z) = -\frac{\sigma V_P}{\theta} = -\frac{\sigma}{\theta}(2\alpha_1(\theta, z)P + \alpha_2(\theta, z)). \quad (20)$$

Before we proceed, we note certain properties regarding the curvature of the maxmin value function $V(P_0, \theta, z) = \alpha_1(\theta, z)P_0^2 + \alpha_2(\theta, z)P_0 + \alpha_3(\theta, z)$ that will be useful later on.

Lemma 1 *The maxmin value function $V(P; \theta, z)$ is*

(a) *Strictly increasing and concave in θ .*

(b) *Strictly decreasing and convex in z . Moreover, the partial derivative V_z is increasing in θ .*

Proof. Part (a) can be established either through cumbersome differentiation, or by referring to Section 5.2 of Hansen et al. [11] and noting that, in our case, Assumption 5.5 holds.

We now turn to part (b). Differentiating $\alpha_1(\theta, z)$ with respect to z , yields

$$\frac{\partial}{\partial z}\alpha_1(\theta, z) = \frac{-s}{\sqrt{(2m + \rho)^2 + 8sz(1 - \frac{\sigma^2}{\theta})}} < 0 \quad (21)$$

which is clearly increasing in θ and z . Doing the same for $\alpha_2(\theta, z)$ we obtain

$$\frac{\partial}{\partial z}\alpha_2(\theta, z) = \frac{2A(\rho + m)\frac{\partial}{\partial z}\alpha_1(\theta, z)}{\left(2\alpha_1(\theta, z)(\frac{\sigma^2}{\theta} - 1) + \rho + m\right)^2} = \frac{2A(\rho + m)\frac{\partial}{\partial z}\alpha_1(\theta, z)}{\left(\frac{\rho}{2} + \frac{1}{2}\sqrt{(2m + \rho)^2 + 8sz(1 - \frac{\sigma^2}{\theta})}\right)^2} < 0, \quad (22)$$

⁵To get rid of uninteresting mathematical complications, we only consider $\theta > \sigma^2$.

which is, therefore, also increasing in θ and z . Turning to $\alpha_3(\theta, z)$

$$\begin{aligned}\frac{\partial}{\partial z}\alpha_3(\theta, z) &= \frac{\sigma^2}{\rho} \frac{\partial}{\partial z}\alpha_1(\theta, z) + \frac{\frac{\partial}{\partial z}\alpha_2(\theta, z)}{\rho} \left[A + \alpha_2(\theta, z) \left(1 - \frac{\sigma^2}{\theta} \right) \right] \\ &= \frac{1}{\rho} \left[\sigma^2 \frac{\partial}{\partial z}\alpha_1(\theta, z) + \frac{\partial}{\partial z}\alpha_2(\theta, z) \frac{2A(m + \rho)}{\rho + \sqrt{(2m + \rho)^2 + 8sz(1 - \frac{\sigma^2}{\theta})}} \right] < 0,\end{aligned}\quad (23)$$

so we see that this too is negative and increasing in θ and z . ■

Eqs.(21) (22),and (23) further suggest that $\frac{\partial}{\partial z}V(\theta, z, P_0)$ does not diverge at $z = 0$ so that

$$\lim_{z \rightarrow 0} \frac{\partial}{\partial z}V(\theta, z, P_0) > -\infty. \quad (24)$$

It is easy to see that $\alpha_1(\theta, z)$ and $\alpha_2(\theta, z)$ are negative and increasing in θ . Thus, Eq. (19) directly suggests the presence of a precautionary principle in emission mitigation: *the more uncertainty exists over pollution dynamics, the more one chooses to mitigate emissions at a given pollution level P* . It is equally straightforward to notice that $\alpha_1(\theta, z)$ and $\alpha_2(\theta, z)$ are decreasing in z . Thus, given a fixed level of misspecification θ , the less we invest in damage control technology, the more we mitigate emissions. These results are not at all surprising.

3.2 Characterizing the worst-case pollution accumulation process

Eq. (20) specifies the worst-case misspecification of our model, given θ . Substituting it into our robust pollution dynamics (7) yields

$$dP(t) = \left(\underbrace{E - \frac{\sigma^2}{\theta}\alpha_2(\theta, z)}_{\text{Effect 1}} - \left[m + \underbrace{\frac{2\sigma^2}{\theta}\alpha_1(\theta, z)}_{\text{Effect 2}} \right] P(t) \right) dt + \sigma dB(t) \quad (25)$$

Eq. (25) points to two pernicious effects of model misspecification. First, there now exists a constant positive drift term (Effect 1) equal to

$$-\frac{\sigma^2}{\theta}\alpha_2(\theta, z) > 0.$$

Second, the environment's self-cleaning capacity has been reduced (Effect 2) by an amount

$$\frac{2\sigma^2}{\theta}\alpha_1(\theta, z) < 0.$$

As we saw earlier, the government reacts to this worst-case scenario by adopting an emissions strategy E^* given by Eq. (19). Thus, at optimality, the worst-case pollution process, call it P^* , is governed by the following stochastic differential equation

$$dP^*(t) = (E^* - mP^*(t) + \sigma \cdot v^*(t))dt + \sigma dB(t), \quad (26)$$

which, given Eqs. (20) and (19), reduces to

$$dP^*(t) = -\left[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m\right] \left(\frac{A + \alpha_2(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right)}{-[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m]} - P^*(t)\right) dt + \sigma dB(t) \quad (27)$$

where we note that

$$\begin{aligned} -\left[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m\right] &= \frac{\sqrt{(2m + \rho)^2 + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)} - \rho}{2} \\ A + \alpha_2(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) &= \frac{2A(m + \rho)}{\rho + \sqrt{(2m + \rho)^2 + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)}} \\ \Rightarrow \frac{A + \alpha_2(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right)}{-\left[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m\right]} &= \frac{4A(m + \rho)}{4m^2 + 4m\rho + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)} \end{aligned}$$

SDE (27) is an instance of the well-known Ornstein-Uhlenbeck process so that we may deduce the following:

Proposition 2 *Stochastic differential equation (27) has a unique solution given by a Gaussian diffusion process $\{P^*(\theta, z, t) : t \geq 0\}$ where*

(a) $P^*(\theta, z, t)$ has expectation

$$\mathbf{E}[P^*(\theta, z, t)] = \hat{P}_0 e^{[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m]t} + \frac{4A(m + \rho)}{4m^2 + 4m\rho + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)} \left[1 - e^{[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m]t}\right],$$

and variance

$$\mathbf{Var}[\hat{P}^*(\theta, z, t)] = \frac{\sigma^2}{\sqrt{(2m + \rho)^2 + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)} - \rho} \left[1 - e^{2[2\alpha_1(\theta, z)\left(1 - \frac{\sigma^2}{\theta}\right) - m]t}\right]^2.$$

(b) $\{P^*(\theta, z, t) : t \geq 0\}$ has a stationary distribution that is $N\left(\frac{4A(m + \rho)}{4m^2 + 4m\rho + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)}, \frac{\sigma^2}{\sqrt{(2m + \rho)^2 + 8sz\left(1 - \frac{\sigma^2}{\theta}\right)} - \rho}\right)$.

Proposition 2 also agrees with our intuition. In steady state, the expected value and variance of the worst-case pollution levels are decreasing in θ and z .

Given Proposition 2 and the explicit characterization of the first and second moments of $P^*(\theta, z, t)$, we may also note that the worst-case entropy of our misspecified model is equal to:

$$\begin{aligned} R(Q^*(\theta, z)) &= \int_0^\infty e^{-\rho t} \frac{1}{2} \mathbf{E}_{Q^*}[v^*(t)^2] dt \\ &= \frac{\sigma^2}{2\theta^2} \int_0^\infty e^{-\rho t} \left[4\alpha_1^2(\theta, z) \left(\mathbf{E}[P^*(\theta, z, t)]^2 + \mathbf{Var}[\hat{P}^*(\theta, z, t)]\right) \right. \\ &\quad \left. + 4\alpha_1(\theta, z)\alpha_2(\theta, z)\mathbf{E}[P^*(\theta, z, t)] + \alpha_2^2(\theta, z)\right] dt. \end{aligned} \quad (28)$$

Thus, we are able to (via Proposition 1) directly associate an entropy bound η to a given ambiguity parameter θ , such that the respective multiplier (10) and constraint (11) robust control problems admit identical solutions.⁶

4 Solving the optimal investment problem

Suppose that at time 0 a policy maker wants to decide how much to invest in damage-control technology. In our notation, he or she would like to choose a value of z . Statistical evidence and climate science suggests a possible model misspecification for the pollution accumulation dynamics that is captured through an ambiguity parameter θ . The policy maker takes this misspecification seriously and wishes to guard against it, so that a maxmin criterion is adopted over future welfare. Recall that $V(P_0, \theta, z)$ denotes the maxmin value of a constraint problem multiplier θ with technology adoption z , at initial pollution P_0 . Thus, at time 0, the policy maker wishes to solve the following optimization problem

$$\max_{z \in [0,1]} V(P_0, \theta, z) - \phi(z) \quad (29)$$

We now make the following assumption on our model primitives, in order to ensure a “single-crossing” property on the objective function of (29).

Assumption 1 *Parameters s, m, ρ, A , initial pollution levels P_0 , and cost function ϕ satisfy*

$$\lim_{\theta \rightarrow \infty} \frac{\partial}{\partial z} V(\theta, 1, P_0) < \phi'(1). \quad (30)$$

Assumption 1 ensures that the optimization problem (29) corresponding to the case of pure risk (and no uncertainty) has a unique solution. In the rest of the paper we suppose that Assumption 1 holds.

Lemma 2 *Consider optimization problem (29). There exists a unique optimal level of damage-control investment z , call it $z^*(\theta)$, that satisfies*

$$\begin{aligned} \frac{\partial}{\partial z} V(P_0, \theta, z) &> \phi'(z), \quad \text{for all } z \in [0, z^*(\theta)) \\ \frac{\partial}{\partial z} V(P_0, \theta, z^*(\theta)) &= \phi'(z^*(\theta)), \\ \frac{\partial}{\partial z} V(P_0, \theta, z) &< \phi'(z), \quad \text{for all } z \in (z^*(\theta), 1] \end{aligned}$$

Proof. Recall that ϕ is strictly decreasing and convex, and satisfies $\phi'(0) = -\infty$. This fact, in combination with Lemma 1, Eq. (24), and Assumption 1, establishes the result. ■

⁶On a technical note straightforward, if cumbersome, algebra verifies that $\frac{d}{d\theta} R(Q^*(\theta, z)) < 0$.

Having established this result, we are now ready to investigate the properties of the optimal solution of (29). We first address optimal investment in damage-control technology and prove that it, indeed, is consistent to the PP.

Theorem 1 *Optimal damage-control investment strictly increases in model uncertainty. In other words, $z^*(\theta)$ is strictly increasing in θ .*

Proof. Consider $\theta_2 > \theta_1$ and the associated optimal investment decisions $z^*(\theta_1), z^*(\theta_2)$. Lemma 2 implies that $z^*(\theta_1)$ uniquely satisfies

$$\frac{\partial}{\partial z} V(\theta_1, z^*(\theta_1), P_0) = \phi'(z^*(\theta_1)).$$

Now, Lemma 1 implies that

$$\frac{\partial}{\partial z} V(\theta_2, z^*(\theta_1), P_0) > \frac{\partial}{\partial z} V(\theta_1, z^*(\theta_1), P_0) = \phi'(z^*(\theta_1)).$$

Consequently, Lemma 2 implies that

$$\frac{\partial}{\partial z} V(\theta_2, z, P_0) > \phi'(z), \quad \text{for all } z \in [0, z^*(\theta_1)],$$

so that it must be the case that $z^*(\theta_2) > z^*(\theta_1)$. ■

Theorem 1 confirms the PP in the case of damage control investment. We now address this question in the context of optimal mitigation policies.

Lemma 3 *The optimal solution of optimization problem (29) is such that the values of $\frac{d}{d\theta}\alpha_1(\theta, z^*(\theta))$ and $\frac{d}{d\theta}\alpha_2(\theta, z^*(\theta))$ can be positive or negative. In particular*

(a)

$$\frac{d}{d\theta}\alpha_1(\theta, z^*(\theta)) < 0 \Leftrightarrow \frac{dz^*}{d\theta}(\theta) > \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta))}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))}$$

(b)

$$\frac{d}{d\theta}\alpha_2(\theta, z^*(\theta)) < 0 \Leftrightarrow \frac{dz^*}{d\theta}(\theta) > \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta)) - \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))}$$

Proof. Consider θ and the associated optimal $z^*(\theta)$. We begin with part (a) and consider $\alpha_1(\theta, z^*(\theta))$. The result immediately follows from differentiating with respect to θ and recalling the negative sign of $\frac{\partial\alpha_1}{\partial z}(\theta, z)$:

$$\frac{d}{d\theta}\alpha_1(\theta, z^*(\theta)) = \frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta)) + \frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))\frac{dz^*}{d\theta}(\theta). \quad (31)$$

Moving on to part (b), we first refer to Eq. (17). Straightforward differentiation establishes that

$$\frac{d}{d\theta}\alpha_2(\theta, z^*(\theta)) = \frac{2A\left(\frac{d}{d\theta}\alpha_1(\theta, z^*(\theta))(\rho + m) + \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2}\right)}{\left(2\alpha_1(\theta, z^*(\theta))\left(\frac{\sigma^2}{\theta} - 1\right) + \rho + m\right)^2} \quad (32)$$

Part (b) now follows from Eqs. (31) and (32). ■

Theorem 2 Consider a neighborhood of θ , say $[\theta_{\min}, \theta_{\max}]$. If $z^*(\theta)$ satisfies

(a)

$$\frac{dz^*}{d\theta}(\theta) > \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta)) - \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in [\theta_{\min}, \theta_{\max}], \quad (33)$$

then emissions (mitigation) are unambiguously decreasing (increasing) in θ in $[\theta_{\min}, \theta_{\max}]$;

(b)

$$\frac{dz^*}{d\theta}(\theta) < \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta))}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in [\theta_{\min}, \theta_{\max}], \quad (34)$$

then emissions (mitigation) are unambiguously increasing (decreasing) in θ in $[\theta_{\min}, \theta_{\max}]$;

(c)

$$\frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta))}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))} < \frac{dz^*}{d\theta}(\theta) < \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta)) - \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in [\theta_{\min}, \theta_{\max}] \quad (35)$$

then emissions will be decreasing in θ for $\theta \in [\theta_{\min}, \theta_{\max}]$ if and only if current pollution levels are high enough.

Proof. Follows immediately from Lemma 3 and the fact that, as Eq. (19) suggests, $E^*(\theta, z, P) = A + 2\alpha_1(\theta, z)P + \alpha_2(\theta, z)$. ■

Remarks. From Theorem 1 we know that $\frac{dz^*}{d\theta}(\theta) > 0$. Moreover, straightforward, if cumbersome, algebra establishes that

$$\frac{\partial\alpha_1}{\partial\theta}(\theta, z) \geq 0, \quad \frac{\partial\alpha_1}{\partial\theta}(\theta, 0) = 0 \quad (36)$$

and Eq. (21) shows that

$$\frac{\partial\alpha_1}{\partial z}(\theta, z) < 0, \quad \frac{\partial^2\alpha_1}{\partial\theta\partial z}(\theta, z) > 0, \quad \frac{\partial\alpha_1}{\partial z}(\theta, 0) = -\frac{s}{2m + \rho} \quad (37)$$

Therefore, the conditions of Theorem 2 are not generically false so that it is, theoretically, possible for mitigation to down up as uncertainty goes up. Moreover, Eqs. (16), (36), and (37) imply that the right-hand-side of Eq. (33) is increasing in $z^*(\theta)$ and satisfies

$$\lim_{z^*(\theta) \rightarrow 0} \frac{-\frac{\partial\alpha_1}{\partial\theta}(\theta, z^*(\theta)) - \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial\alpha_1}{\partial z}(\theta, z^*(\theta))} = 0.$$

The above discussion suggests that, in some instances, the precautionary principle may only be manifest in our choice of damage-control technology, particularly in cases where z^* is low and sensitive to changes in θ . Hence, we arrive at the following corollary of Theorem 2.

Corollary 1 *The right-hand side of Eq. (33) is increasing in $z^*(\theta)$, and vanishes at $z^*(\theta) = 0$. Hence, for high enough levels of optimal damage-control investment (i.e., low enough $z^*(\theta)$), emissions will be decreasing in θ , provided the rate of change of $z^*(\theta)$ is high enough. In other words, if optimal levels of damage-control investment are both high enough and sufficiently sensitive to changes in uncertainty, then we observe a reversal of the PP with regard to mitigation.*

The intuition behind this result can be described in the following way: If damage-control investment is sensitive to θ , then an increase in uncertainty will cause a large increase in damage-control investment, which in turn will reduce damages from time 0 onwards. If this reduction is sufficiently large then, since more mitigation is also costly, incentives to mitigate weaken to the extent that mitigation is actually reduced. In this case we observe that when uncertainty increases, damage-control investment and mitigation become *substitutes* rather than *complements*.

A heuristic algorithm for solving the constraint robust control problem. Clearly, the above analysis captures model misspecification through a multiplier robust control problem. However, as we mentioned earlier, an ambiguity parameter θ tends to have no direct physical meaning; instead, it functions as a Lagrange multiplier of a constraint on relative entropy. By contrast, relative entropy bounds *do* have physical meaning as their determination is based on physical principles and statistical analysis of time series data. Thus, the question naturally arises: Can we derive optimal mitigation and damage-control policies under a constraint η on relative entropy? The answer is most likely yes, through the following heuristic “algorithm”:

1. Suppose that physical science and time series analysis suggests a model misspecification of η^* . Pick $\theta > \sigma^2$ at random.
2. Solve optimization problem (29) and determine $z^*(\theta)$.
3. Use Eq. (28) to calculate $\eta = R(Q^*(\theta, z^*(\theta)))$. By Proposition 1 the constraint problem with an entropy bound η will admit an identical solution as the one we have just computed through the multiplier problem θ .
4. If $\eta > \eta^*$, pick a bigger θ and repeat steps 2 and 3. Otherwise, pick a smaller θ and do the same. Repeat until the derived η is close enough to the given η^* .

Why is the above algorithm heuristic? It is because we have not formally proved that

$$\frac{d}{d\theta} R(Q^*(\theta, z^*(\theta))) < 0,$$

so that step 4 may not necessarily lead to convergence. Addressing the monotonicity of $R(Q^*(\theta, z^*(\theta)))$ with respect to θ in full generality is complicated by the countervailing effects that θ and z have on entropy. In particular, recall from Theorem 1 that z^* is increasing in θ . Moreover, note from Eq. (28) that

$$\frac{\partial}{\partial \theta} R < 0, \quad \frac{\partial}{\partial z} R > 0.$$

Thus, it becomes clear that we cannot immediately deduce that worst-case entropy will be decreasing in θ . Still, numerical results indicate that an increase in θ significantly outweighs its attendant increase of z^* (see Section 5 and Table 1), and we are confident that the above algorithm is a useful way of embedding entropy bounds in our robust control problem.

5 The effect of damage-control cost ϕ on precaution

In this section we explore how changes in the cost function ϕ affect optimal mitigation and damage control decisions. We focus on the following family of cost functions that is consistent with our model assumptions

$$\phi(z; k) = \frac{1}{z^k} - 1, \quad k \geq 1, \tag{38}$$

so that $\phi(z; k_1) > \phi(z; k_2)$ (unless of course $z = 0$ or 1) and $\phi'(z; k_1) < \phi'(z; k_2)$ whenever $k_1 > k_2$. Hence cost (marginal cost) is increasing (decreasing) in k . We begin with a natural result.

Proposition 3 *Fix a level of uncertainty θ and consider a family of optimization problems (29), parametrized according to Eq. (38).*

- (a) *Optimal values of $z^*(\theta; k)$ are increasing in k . In other words, optimal levels of damage-control investment are decreasing in the cost of damage control technology.*
- (b) *Optimal emissions are decreasing in k . In other words, optimal levels of mitigation are increasing in the cost of damage-control technology.*

Proof. Part (a) follows from Lemma 2 and the fact that $\phi'(z; k_1) < \phi'(z; k_2)$ whenever $k_2 > k_1$. Part (b) follows from part (a) and Eqs. (16) and (17). ■

Proposition 3 is not surprising. The more expensive damage-control technology is, the less we can expect to invest in it. Moreover, this decrease in damage control means that additional mitigation is necessary, to protect against high pollution concentrations. Proposition 3 holds for every level of uncertainty $\theta > \sigma^2$.

A numerical example. We now examine the effect of ϕ on precautionary behavior, as we vary θ . To make the analysis concrete, we focus on the following problem instance:

$$P_0 = 100, \quad m = \rho = .03, \quad A = 100, \quad s = 1, \quad \sigma^2 = 1.$$




Figure 1: $\frac{z^*(\theta;k)}{z^*(\infty;k)}$ as a function of θ for $k = 1, 2, \dots, 6$.

We already know from Theorem 1 that optimal damage-control investment will be increasing in uncertainty, i.e., that $z^*(\theta; k)$ is increasing in θ , for all cost functions (38). Indeed this can be seen in Figure 1. On the other hand, Figure 2 illustrates Proposition 3 and shows how, given a level of uncertainty θ , optimal damage-control investment will be decreasing in the cost of technology.

But while our choice of cost function does not affect the precautionary principle vis-a-vis optimal damage-control investment, the situation is not so simple in the case of mitigation. Indeed, all three cases of Theorem 2 may occur for different values of k . The case in which $k = 1.5$ corresponds to part (a) of Theorem 2 (see Figure 3), $k = 2$ to part (c) (see Figure 4), and $k = 3.5$ (see Figure 5) to part (b).

Taken together, our numerical results indicate that when the cost of damage-control technology is low enough (i.e., when k is small) we are in a situation in which the precautionary principle with regard to mitigation unambiguously does not hold. This can be explained by the fact that in such cases, the government will choose to invest a lot in damage control technology and the intuition of Corollary 1 will be applicable.

6 Conclusion

The present paper analyzed optimal pollution control policy under Knightian uncertainty by adopting the robust control framework of Hansen and Sargent [10]. Allowing for a one-time investment in




Figure 2: $z^*(\theta; k)$ as a function of θ for $k = 1, 2, \dots, 6$.




Figure 3: Emissions as a function of P for different θ : ($k = 1.5, P_0 = 100$). This case corresponds to part (a) of Theorem 2.




Figure 4: Emissions as a function of P for different θ : ($k = 2, P_0 = 100$). This case corresponds to part (c) of Theorem 2.




Figure 5: Emissions as a function of P for different θ : ($k = 3.5, P_0 = 100$). This case corresponds to part (b) of Theorem 2.

k	θ	$\mathbf{E}[P^*]$	$\mathbf{Var}[P^*]$	η
1.5	1.1	3112	15.9	37131
1.5	1.5	2584	14.1	18663
1.5	2	2267	13	9995
1.5	3	1967	11.8	4221
1.5	4	1814	11.2	2312
2	1.1	2592	14.1	456998
2	1.5	1346	9.4	166074
2	2	874	7.2	70910
2	3	551	5.4	24063
2	4	400	4.7	11841
3.5	1.1	720	6.45	8333650
3.5	1.5	119	2.38	614720
3.5	2	62.8	1.7	163497
3.5	3	35.8	1.2	41990
3.5	4	7.93	1.04	18818

Table 1: Steady-state expectation/variance and relative entropy of worst-case pollution levels P^* , as a function of k and θ .

damage-control technology, in addition to gradual emissions mitigation, we studied the applicability of a precautionary principle with respect to *both* damage control and mitigation. Our main finding is that while investment in damage-control technology is always increasing in uncertainty, optimal mitigation is not. Indeed, if optimal levels of damage-control investment are both high enough and sufficiently sensitive to changes in uncertainty, then we observe a reversal of the PP with regard to mitigation. We consider this to be an interesting consequence of the interaction between two different ways of reducing pollution-related damages.

From a normative standpoint our analysis implies that, depending on the cost of damage-control technology and the magnitude of uncertainty, it may be preferable to be precautionary now by undertaking large damage-control investment, and not be particularly precautionary with respect to future mitigation policy. Indeed, when this is the case, current damage-control investment and future mitigation act as substitutes. On the other hand, when damage-control investment is costly, it can act as a complement to future mitigation and an increase in uncertainty induces precaution with respect to both policy actions. We think that these results provide new insights into the question of when and how to be precautionary in the design of pollution control policy in the context of robust decision making.⁷ Interesting future work along similar lines would extend the basic model to incorporate dynamic damage-control investment, more intricate pollution dynamics, as well as lower bounds on emissions that would reflect concerns about irreversibility.

References

- [1] T. Asano (2010), “Precautionary Principle and the Optimal Timing of Environmental Policy under Ambiguity,” *Environmental and Resource Economics*, 47, 173–196
- [2] Z. Chen and L. Epstein (2002), “Ambiguity, Risk, and Asset Returns in Continuous Time,” *Econometrica*, 70, 1403–1443.
- [3] E. Dockner and N. Van Long (1993), “International Pollution Control: Cooperative vs. Noncooperative Strategies,” *Journal of Environmental Economics and Management*, 24, 13–29.
- [4] D. Ellsberg (1961), “Risk Ambiguity and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- [5] L. Epstein and M. Schneider (2003), “Recursive Multiple Priors,” *Journal of Economic Theory*, 113, 1–31.
- [6] W. Fleming and P. Souganidis (1989), “On the Existence of Value Function of Two-Player, Zero Sum Stochastic Differential Games,” *Indiana University Mathematics Journal*, 3, 293–314.
- [7] I. Gilboa and D. Schmeidler (1989), “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- [8] C. Gollier, J. Gierlinger (2008), “Socially efficient discounting under ambiguity aversion,” Working paper.

⁷In a sense, our analysis might be interpreted as supporting the argument for taking strong immediate action with regard to climate policy under deep uncertainty (e.g. Weitzman [26]).

- [9] F. Gonzalez (2008), “Precautionary Principle and Robustness for a Stock Pollutant with Multiplicative Risk,” *Environmental and Resource Economics*, 41, 25–46.
- [10] L. P. Hansen and T. Sargent (2001), “Robust Control and Model Uncertainty,” *American Economic Review P & P*, 91, 60-66.
- [11] L. P. Hansen, T. Sargent, G. Turhumambetova, and N. Williams (2006), “Robust Control and Model Misspecification,” *Journal of Economic Theory*, 128, 45-90.
- [12] M. Hoel and L. Karp (2001), “Taxes Versus Quotas for a Stock Pollutant under Multiplicative Uncertainty,” *Journal of Public Economics*, 82, 91–114.
- [13] L. Karp (2009), “Sacrifice, discounting and climate policy: Five questions,” Working paper.
- [14] P. Klibanoff, M. Marinacci, S. Mukerji (2005), “A Smooth Model of Decision Making Under Ambiguity,” *Econometrica*, 73, 1849–1892
- [15] F. Knight (1921), *Risk, Uncertainty, and Profit* Houghton Mifflin, USA.
- [16] A. Millner, S. Dietz, and G. Heal (2010) “Ambiguity and climate policy,” Working paper.
- [17] J. von Neumann and O. Morgenstern. (1944), *Theory of Games and Economic Behavior*, Princeton University Press.
- [18] W. Nordhaus (2008), *A Question of Balance: Economic Modeling of Global Warming*, Yale University Press.
- [19] W. Nordhaus (2009), “An Analysis of the Dismal Theorem,” working paper.
- [20] C. Roseta-Palma and A. Xepapadeas (2004) “Robust control in water management,” *Journal of Risk and Uncertainty* 29, 21–34.
- [21] L. J. Savage (1954), *The Foundations of Statistics*, Wiley, New York.
- [22] N. Stern (2007), *Stern Review: The Economics of Climate Change*, Cambridge University Press
- [23] Sterner, T., Persson, M. (2008). “An even sterner review: Introducing relative prices into the discounting debate,” *Review of Environmental Economics and Policy* 2, 61-76.
- [24] Traeger, C.P. (2009). Recent developments in the intertemporal modeling of uncertainty. *Annual Review of Resource Economics* 1, 261-286.
- [25] G. Vardas and A. Xepapadeas (2010), “Model Uncertainty, Ambiguity, and the Precautionary Principle: Implications for Biodiversity Management,” *Environmental and Resource Economics*, 45, 379–404.
- [26] M. Weitzman (2009), “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*, 91, 1–19
- [27] M. Weitzman (2010), “What is the “damages function” for global warming -and what difference might it make?” *Climate Change Economics* 1, 57-69.