Pollution control with uncertain stock dynamics: when, and how, to be precautious

Stergios Athanassoglou[*], Anastasios Xepapadeas[†]

February 2011; revised August 2011

[*]Corresponding author. Fondazione Eni Enrico Mattei and Euro-Mediterranean Center for Climate Change. Corso Magenta 63, 20123 Milan, Italy. +390252036983, athanassoglou@gmail.com.

[†]Athens University of Economics and Business and Beijer Fellow. 76 Patission Street, 104 34 Athens, Greece. xepapad@aueb.gr. We are grateful to the Editor and two anonymous referees for their comments on an earlier version of the paper. Their critical remarks and suggestions led to many substantial improvements in both form and content. We thank Valentina Bosetti, Ben Groom and seminar participants at the Fondazione Eni Enrico Mattei, University of Brescia, and 2011 EAERE Conference for useful comments.

Abstract

The precautionary principle (PP) applied to environmental policy stipulates that, in the presence of uncertainty, society must 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 with uncertain (in the Knightian sense) stock dynamics, using the robust control framework of Hansen and Sargent [12]. Optimal investment in damage control is found to be increasing in the degree of uncertainty, thus confirming the conventional PP wisdom. Optimal mitigation decisions, however, need not always comport with the PP. In particular, when damage-control investment is both sufficiently cheap and sensitive to changes in uncertainty, damage-control investment and mitigation may act as substitutes and a PP with respect to the latter can be unambiguously irrational. The theoretical results are consequently applied to a linear-quadratic model of climate change calibrated by Karp and Zhang [20]. The analysis suggests that a reversal of the PP with respect to mitigation, while theoretically possible, is very unlikely.

Keywords: Knightian uncertainty, robust control, precautionary principle, pollution control, stock dynamics

JEL classifications: C61, D80, D81

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. There are many compelling reasons for its primacy: expected utility theory has solid theoretical underpinnings, going back to the work of von Neumann and Morgenstern [26] and Savage [29], is conceptually intuitive, 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 [34], and (b) the high sensitivity of model outputs to controversial modeling assumptions (for instance, the functional form of the chosen damage function [31, 35] and the value of the social discount rate). As a result, separate models may arrive at dramatically different policy recommendations, generating significant uncertainty over the magnitude and timing of desirable policy.[1]

A general guide for crafting policy under such uncertain conditions can be found in the formulation of a *precautionary principle* (PP). In plain English, the PP basically codifies the age-old mantra "better safe than sorry". Here is the way it was expressed as Principle 15 of the Rio Declaration, in the context of the 1992 United Nations Earth Summit:

> "In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation."[2]

The Wingspread Statement, formulated at the 1998 Wingspread Conference on the Precautionary Principle, goes even further:

[1] William Nordhaus' DICE model [27] and the Stern Report [30] are the canonical examples of this deep divergence within the context of climate change economics.

[2] http://www.unep.org/Documents.multilingual/Default.asp?DocumentID=78&ArticleID=1163

> "When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically."[3]

In our work we focus on yet another variation still, one which involves the adaptation of policy to *changing* levels of uncertainty. In particular, we consider an extension of the PP that prescribes an increase in the stringency of precautionary policy as the degree of uncertainty grows. While this statement does not necessarily follow from either of the above formulations of the PP, we believe it to be a defensible extension of its overarching logic.

To ground our study on a rigorous quantitative basis, we take the broad term of "uncertainty" to mean an inability to posit precise probabilistic structure to physical and economic models. This derives from the concept of uncertainty as introduced by Knight [23] to represent a situation where a decisionmaker lacks adequate information to assign probabilities to events. Knight argued that this deeper kind uncertainty is quite common in economic decisionmaking, and thus deserving of systematic study. Knightian uncertainty is contrasted to risk (measurable or probabilistic uncertainty) where probabilistic structure can be fully captured by a *single* Bayesian prior. There is considerable evidence that it may provide a more appropriate modeling framework for many applications in environmental economics, especially climate change [34, 25].

Inspired by the work of Knight 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.[4] Klibanoff et al. [21, 22] developed an axiomatic framework, the "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

[3] http://www.sehn.org/state.html#w

[4] The related decision-theoretic literature is both vast and deep so the following remarks are by no means meant to be exhaustive. We purely focus on the few contributions that are directly relevant for our purposes.

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 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.* A truly compelling and innovative feature of the smooth ambiguity model is that it allows for a separation between ambiguity (the set Π and the second-order distribution p) and a decisionmaker's attitude (i.e., aversion) towards it, nesting in a smooth fashion the entire continuum between simple aggregation of the prior π's (ambiguity neutrality) to absolute focus on the worst-case (absolute ambiguity aversion). Comparative statics exercises involving the above are relatively easy to perform (at least in the static version of the model) and generate rich and insightful results.

In recent years the smooth ambiguity framework has been applied to a number of issues in environmental economics (Gollier and Gierlinger [9], Treich [32], Millner et al. [25], Lemoine and Traeger [24]). However, despite its prominent role in the recent literature, the smooth ambiguity model seems (at least to us) to have more of a positive instead of a normative focus, and questions about how to calibrate agents' ambiguity aversion in environmental settings appear difficult to address. As an example, consider global climate-change policy: it is unclear to us how one could, or even should, use Ellsberg-type thought experiments to calibrate ambiguity aversion parameters on whose ultimate basis normatively-appealing emissions trajectories will be determined. 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. On a final note, it is worth mentioning that the dynamic version of the smooth ambiguity model [22] seems to pose nontrivial tractability challenges, so that (at times) only the utility of very simple, exogenously given, policies can be computed [25].

Our Focus: Robust Control. In a seminal contribution, Gilboa and Schmeidler [8] 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 compelling, to optimize over the worst-case outcome (i.e. the worst-case *prior*) that may conceivably come to pass. 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.

Motivated by the possibility of model misspecification in macroeconomics, Hansen and Sargent [12] and Hansen et al. [15] extended Gilboa and Schmeidler's insight to continuous-time dynamic optimization problems, introducing the concept of robust control to economic environments. They showed how standard dynamic programming techniques can be modified to yield robust solutions to problems in which the underlying stochastic nature of the model is not perfectly known.[5] In their work, the degree of misspecification is a model input, so that decision makers can test the sensitivity of a proposed solution with respect to the model's presumed uncertainty. Lacking complex formal characterizations similar to Klibanoff et al. [21, 22] and Epstein and Schneider [5], the focus of Hansen-Sargent robustness project seems to be as much practical as it is theoretical, if not more.[6]

Finally, we should also 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

[5] In Section 2 we discuss the relationship of robust control to risk-sensitive control theory, developed earlier in the engineering and control literature.

[6] There are, however, some important shortcomings to the robust control framework that bear mentioning. In contrast to the smooth ambiguity model, the max-min setting of robust control cannot disentangle ambiguity and ambiguity attitude (as ambiguity attitude is fixed) and preferences will, in general, be kinked. Moreover, the basic version of the model that we use does not allow for learning over time so that it is assumed that a decisionmaker cannot re-adjust his model misspecification to reflect historical data. Later work by Hansen and Sargent [13] addresses this concern.

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 (it is larger, and therefore more general), and their evolution over time.[7] A recent application of RMP in environmental economics can be found in Asano [1], who studies the optimal timing of environmental policy under ambiguity.

Our contribution. In recent years the Hansen-Sargent framework has slowly begun to make its way into environmental economics. Gonzales [10] applied robust control to the regulation of a stock pollutant under multiplicative uncertainty introduced by Hoel and Karp [16]. Roseta-Palma and Xepapadeas [28] studied water management under ambiguity, while Vardas and Xepapadeas [33] did the same in the context of biodiversity management. Funke and Paetz [7] applied the robust control framework to a numerical model of climate change while Xepapadeas [37] studied an international game of pollution control under cooperative and non-cooperative assumptions on countries' behavior.

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 pollution control, studied by Dockner and Van Long [3] among many others, to allow for (a) misspecification of stock dynamics and (b) the possibility of investment in *damage-control* technology that alleviates the effects of pollutant stock accumulation. In the context of climate change, examples of this kind of damage-control investment can be found in the construction of large-scale civil engineering projects, substantial R & D in geoengineering, and the construction of new urban environments to accommodate potential forced migration. It is distinct from direct emissions *mitigation*, which is traditionally attempted through economic instruments such as taxes, emissions quotas, and assorted command-and-control measures.

We assume the presence of a benevolent government (or, conversely, a group of cooper-

[7]For more details the reader is referred to section 5 in Epstein and Schneider [5] and section 9 in Hansen et al. [15].

ating 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 framework, we introduce Knightian uncertainty into the basic model and study the effect of model misspecification on optimal mitigation and damage-control decisions. We focus on uncertainty surrounding the *pollution stock*, and in particular its accumulation dynamics. Specifically, uncertainty is introduced in the underlying diffusion process, reflecting concerns about our benchmark probabilistic model such as: (a) a miscalculation of exogenous sources of emissions, (b) a miscalculation of the natural pollution decay rate, and (c) an ignorance of more complex dynamic structure involving irreversibility, feedback or hysteresis effects.

In contrast to previous contributions [10, 33, 7, 37] we provide an explicit analytical solution to the maxmin problem that clarifies the structure of robust feedback policies.[8] Moreover, to the best of our knowledge, our paper is the first to (a) completely characterize and physically interpret the stochastic pollution dynamics that result, and (b) attach a statistically meaningful, as well as analytically tractable, parameter (entropy bound) on the degree of model misspecification. These insights prove especially useful in our paper's numerical exercise.

Our primary focus is normative. Ex-ante, one may 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* decrease emissions *and* invest in damage control. Indeed, since higher

[8] To be fair, [10, 7] studied discrete-time models which do not lend themselves to nice closed-form solutions and in which even steady-state results are hard to come by (Gonzalez [10]). Along similar lines, Vardas and Xepapadeas [33] focused on a significantly more complex nonlinear model, which they had to linearize around the steady state in order to derive some insight into the structure of optimal solutions. Xepapadeas [37] studied a similar linear-quadratic model as ours but stopped short of the more complete analysis we perform here, focusing instead on determining the 'cost of preacaution', i.e., the welfare loss that Knightian uncertainty leads to. Finally, Roseta-Palma and Xepapadeas [28] explicitly characterized robust feeback policies for a different model that addressed rainfall uncertainty.

uncertainty translates to the possibility of higher damages from pollutant accumulation, such a finding would not be altogether unreasonable.

However the above conjecture is only partially true. We formally prove 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 agree with the PP and we provide analytical conditions that sway the relationship one way or the other. Initially this result may seem strange; why should we ever emit *more* as uncertainty over damages *increases*? But, upon slightly closer examination, the precautionary result on damage-control technology renders the above not especially surprising or counter-intuitive. The reasoning[9] is simple enough. Keeping uncertainty fixed, emissions are decreasing in damages whilst, keeping damages fixed, they are decreasing in uncertainty. It thus stands to reason why as uncertainty increases and investment in damage-control is ramped up, the net effect on emissions is ambiguous. Indeed, we find that when the cost of damage control is low enough, damage-control investment and mitigation may act as substitutes so that a PP with respect to the latter can be unambiguously irrational.

The theoretical results are consequently applied to a linear-quadratic model of climate change, calibrated by Karp and Zhang [20]. In our simulations we take pains to quantify and carefully calibrate the uncertainty parameter of our model so that our choices reflect realistic cases of model misspecification. Our novel calibration and rigorous interpretation of the numerical results hinge in large part on the theoretical analysis and may be, at least in our view, of independent interest for robust control applications. Our main policy-relevant finding is that emissions can be increasing in uncertainty only when damage-control technology is extremely and most probably unrealistically cheap. Thus, for all practical purposes, when dealing with uncertainty in stock dynamics a precautionary principle with regard to both damage control and mitigation will likely be part of a robust climate-change policy.

[9] elucidated by an anonymous referee

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 the theoretical results with a numerical exercise on a calibrated model of climate change. Section 6 provides concluding remarks.

2 Robust Pollution Control

2.1 Introducing model misspecification and damage control technology

We adopt the standard linear quadratic model of international pollution control analyzed by Dockner and van Long [3], among many others. Output is a function of emissions $F(E)$, 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 following linear differential equation,

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

where $0 < m < 1$ reflects the environment's self cleaning capacity, and $\bar{P} \geq 0$ the pre-industrial level of the pollution stock. Utility is given by $u(F(E)) - D(P)$ where $D(P)$ is a damage function and

$$u(F(E)) = -\frac{b}{2}E^2 + aE, \quad a \geq 0, \;\; b > 0. \tag{2}$$

We modify the standard quadratic damage function $D\left(P\right) = \frac{g}{2}(P - \bar{P})^2, \;\; g > 0$ by allowing for the possibility of investment in damage control (note that damages are identically zero when $P = \bar{P}$). 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

$$D(P,z) = z \cdot \frac{g}{2}(P - \bar{P})^2. \tag{3}$$

Thus, a lower value 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 $\phi(z) : [0,1] \mapsto \Re^+$ that satisfies

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

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

$$\mathrm{d}P(t) = \Big(E - m(P(t) - \bar{P})\Big)\,\mathrm{d}t + \sigma \mathrm{d}B(t), \tag{4}$$

where $\{B(t) : \ t \geq 0\}$ is a Brownian motion on an underlying probabibility space $(\Omega, \mathcal{F}, G)$. Thus, in a world without uncertainty and with fixed damage-control technology, the government's objective is to maximize welfare or

$$\begin{aligned} \max_{E} \qquad & \mathbf{E}\int_0^\infty \mathrm{e}^{-\rho t}\Big[aE - \frac{bE^2}{2} - z\frac{g}{2}(P - \bar{P})^2\Big]\mathrm{d}t \\ \text{subject to:} \qquad & (4), \ \ P(0) = P_0, \end{aligned} \tag{5}$$

where $\rho > 0$ is a discount rate. 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 [12], model misspecification can be reflected by a family of stochastic perturbations to the Brownian motion so that the probabilistic structure implied by stochastic differential equation (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

$$\hat{B}(t) + \int_0^t v(s)\mathrm{d}s, \tag{6}$$

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 mesurable process v could correspond to any number of misspecified or omitted dynamic effects such as: (a) a miscalculation of exogenous sources of emissions, (b) a miscalculation of the natural pollution decay rate, and (c) an ignorance of more complex dynamic structure involving irreversibility, feedback or hysteresis effects. The distortions will be zero when $v \equiv 0$ and the two measures G and Q coincide. Pollution dynamics under model misspecification are given by:

$$\mathrm{d}P(t) = \Big(m\bar{P} + E - mP(t) + \sigma v(t)\Big)\,\mathrm{d}t + \sigma \mathrm{d}B(t). \tag{7}$$

As discussed in Hansen and Sargent [12], the discrepancy between the two measures G and Q is measured through their relative entropy

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

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

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

where $e^{-\rho t}$ is the appropriate discount factor. By modifying the value of η in (9) the decision-maker can control the degree of model misspecification he is willing to consider. In particular, if the decisionmaker can use physical principles and statistical 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 the benchmark pollution dynamics (4) are replaced by Eq. (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

$$\begin{aligned} V(P_0;\theta,z) = \max_E \min_v \quad & \mathbf{E}\int_0^\infty \mathrm{e}^{-\rho t}\Big[aE - \frac{bE^2}{2} - z\cdot\frac{g}{2}(P-\bar{P})^2 + \frac{\theta v^2}{2}\Big]\mathrm{d}t \\ \text{subject to:} \quad & (7),\ \ P(0) = P_0, \end{aligned} \tag{10}$$

while the constraint robust control problem is given by

$$\begin{aligned} V(P_0;\eta,z) = \max_E \min_v \quad & \mathbf{E}\int_0^\infty \mathrm{e}^{-\rho t}\Big[aE - \frac{bE^2}{2} - z\cdot\frac{g}{2}(P-\bar{P})^2\Big]\mathrm{d}t \\ \text{subject to:} \quad & (7),\ \ (9),\ \ P(0) = P_0. \end{aligned} \tag{11}$$

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 missepcification error that the decision-maker is willing to consider. In the multiplier problem (10), the parameter θ can be interpreted as a Lagrangean 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 $\theta \leq \underline{\theta}$, robust control rules cannot be attained. On the other hand when $\theta \to \infty$ 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.[10] In what follows, we primarily focus on the multiplier problem (10) as it is the more analytically tractable problem of the two (Fleming and Souganidis [6]). However, it is

[10]For details the reader is referred to Sections 5 and 7 in Hansen et al. [15].

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 [12], which establishes the following:

Proposition 1 (Prop. 2, Hansen and Sargent [12]) *Suppose V is strictly decreasing in η, $\theta^* \in (\underline{\theta}, +\infty]$, and there exists a solution E^* and v^* (corresponding to measure Q^*) to the multiplier problem (10). Then, E^* also solves the constraint problem (11) for $\eta = \eta^* = R(Q^*)$.*

Relationship to risk-sensitive control. Having defined the multiplier and constraint robust control problems, we briefly comment on the relationship between robust control and earlier research in engineering and applied mathematics. A good deal before Hansen and Sargent's robustness project, control theorists had developed the concept of risk-sensitive control for dynamic optimization problems. Risk-sensitive control theory maximizes a somewhat unconventional objective, namely $-\theta \log \mathbf{E}[\mathrm{e}^{-U/\theta}]$, where U represents an intertemporal utility function and $1/\theta > 0$ a risk-sensitivity parameter. Jacobson [17] and Whittle [36] were the first to show, in a linear-quadratic discrete-time undiscounted finite-horizon model, that the optimal solution to the risk-sensitive problem is identical to the one for the multiplier robust control problem (10) we just discussed. Consequently, James [18] and James and Elliot [19] analyzed continuous-time, nonlinear extensions of the original Jacobson-Whittle model. Hansen and Sargent later [11] extended Jacobson and Whittle's analysis to an infinite-horizon discounted formulation, thus accomodating concerns about time inconsistency of the orginal solutions. For more details on the influence of control theory (risk-sensitive or otherwise) on the economics literature of robust control the reader is referred to section 3.2 of Hansen et al. [15] as well as Hansen and Sargent [14].

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:

$$\rho V = \max_E \min_v \left\{ aE - \frac{bE^2}{2} - z \cdot \frac{g}{2}(P - \bar{P})^2 + \frac{\theta v^2}{2} + V_P(m\bar{P} + E - mP + \sigma v) + \frac{\sigma^2}{2} V_{PP} \right\} \tag{12}$$

Minimizing first with respect to v, we obtain

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

so that Eq. (12) becomes

$$\rho V = \max_E \left\{ aE - \frac{bE^2}{2} - z \cdot \frac{g}{2}(P - \bar{P})^2 + V_P(m\bar{P} + E - mP) + \frac{\sigma^2}{2} V_{PP} - \frac{\sigma^2}{2\theta}(V_P)^2 \right\}. \tag{13}$$

Maximizing with respect to E, we have

$$E^* = \frac{a + V_P}{b}$$

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

$$\rho V = a\frac{a + V_P}{b} - \frac{b}{2}\frac{(a + V_P)^2}{b^2} - z \cdot \frac{g}{2}(P - \bar{P})^2 + V_P\left(m\bar{P} + \frac{a + V_P}{b} - mP\right) + \frac{\sigma^2}{2} V_{PP} - \frac{\sigma^2}{2\theta}(V_P)^2. \tag{14}$$

Straightforward algebra shows that the value function satisfying (14) admits the following simple quadratic form

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

where

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

$$\alpha_2(\theta,z) = \frac{2a\frac{\alpha_1(\theta,z)}{b}+\bar{P}(g+2m\alpha_1(\theta,z))}{2\alpha_1(\theta,z)(\frac{\sigma^2}{\theta}-\frac{1}{b})+\rho+m} \leq 0 \tag{17}$$

$$\alpha_3(\theta,z) = \frac{1}{\rho}\Big[\frac{a^2}{2b}-\frac{zg}{2}\bar{P}^2+\sigma^2\alpha_1(\theta,z)+\alpha_2(\theta,z)(\frac{a}{b}+m\bar{P})+\frac{\alpha_2(\theta,z)^2}{2}\Big(\frac{1}{b}-\frac{\sigma^2}{\theta}\Big)\Big] \tag{18}$$

The value function is well-defined for $\theta > b\sigma^2$ and diverges for $\theta = b\sigma^2$. Hence the Hansen-Sargent breakpoint is equal to $\underline{\theta} = \sigma^2 b$ and we from now on only consider

$$\theta > \underline{\theta} = \sigma^2 b.$$

Max-min optimal emissions E^* satisfy

$$E^*(P,\theta,z) = \frac{a+V_P}{b} = \frac{1}{b}\Big[a+\alpha_2(\theta,z)+2\alpha_1(\theta,z)P\Big], \tag{19}$$

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

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

Before we proceed, we note certain properties regarding the curvature of the maxmin value function $V(P,\theta,z) = \alpha_1(\theta,z)P^2+\alpha_2(\theta,z)P+\alpha_3(\theta,z)$ that will be useful later on. First of all, we re-write the value function in the following way:

$$V(P,\theta,z) = \beta_1(\theta,z)(P-\bar{P})^2+\beta_2(\theta,z)(P-\bar{P})+\beta_3(\theta,z), \tag{21}$$

where simple algebra implies

$$\begin{aligned}\beta_1(\theta,z) &= \alpha_1(\theta,z)\\ \beta_2(\theta,z) &= \alpha_2(\theta,z)+2\alpha_1(\theta,z)\bar{P}\\ \beta_3(\theta,z) &= \alpha_3(\theta,z)+\alpha_1(\theta,z)\bar{P}^2+\alpha_2(\theta,z)\bar{P}.\end{aligned} \tag{22}$$

Lemma 1 *Consider the restricted domain $P \geq \bar{P}$. The maxmin value function $V(P;\theta,z)$ given by Eq. (21) 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 establishished either through differentiation, or by referring to Section 5.2 of Hansen et al. [15] and noting that, in our case, Assumption 5.5 holds.

We now turn to part (b). Let $\Delta(\theta, z) = \sqrt{(2m+\rho)^2 + 4gz(\frac{1}{b} - \frac{\sigma^2}{\theta})}$. Differentiating $\beta_1(\theta, z)$ with respect to z, yields

$$\frac{\partial}{\partial z}\beta_1(\theta, z) = \frac{-g}{2\Delta(\theta, z)} < 0 \tag{23}$$

which is clearly increasing in θ and z. Doing the same for $\beta_2(\theta, z)$ we obtain

$$\frac{\partial}{\partial z}\beta_2(\theta, z) = \frac{-4ag(\rho + m)}{b\Delta(\theta, z)\Big(\rho + \Delta(\theta, z)\Big)^2} < 0, \tag{24}$$

which is also increasing in θ and z. Turning to $\beta_3(\theta, z)$, we obtain

$$\frac{\partial}{\partial z}\beta_3(\theta, z) = -2g\frac{4a^2(m+\rho)^2\theta + (bg\sigma^2 z + m\rho\theta + m^2\theta)(3\rho + \Delta(\theta, z))(\theta - b\sigma^2) + b^2\sigma^2\theta(\rho^3 + \rho^2\Delta(\theta, z))}{b^2\rho\Delta(\theta, z)(\rho + \Delta(\theta, z))^3} \tag{25}$$

so, recalling that we only consider $\theta > b\sigma^2$, we see that this too is negative. Cumbersome differentiation, which can be found in the Appendix, establishes that $\frac{\partial}{\partial z}\beta_3(\theta, z)$ is increasing in θ and z. ■

Eqs.(23), (24), and (25) establish that $\frac{\partial}{\partial z}V(P, \theta, z)$ does not diverge at $z = 0$ so that

$$\lim_{z \to 0} \frac{\partial}{\partial z}V(P, \theta, z) > -\infty. \tag{26}$$

Moreover, it is easy to see that $\beta_1(\theta, z)$ and $\beta_2(\theta, z)$ are negative and increasing in θ. Recalling that

$$E^*(P, \theta, z) = \frac{a + V_P}{b} = \frac{1}{b}\Big[a + \beta_2(\theta, z) + 2\beta_1(\theta, z)(P - \bar{P})\Big], \tag{27}$$

directly suggests the presence of a precautionary principle in emissions mitigation: *the greater the uncertainty over pollution dynamics, the less one chooses to emit at any given pollution level* $P \geq \bar{P}$. Moreover, given a fixed level of misspecification θ, Eq. (27) and the proof of Lemma 1 establish that a similar result applies (i.e., emissions go down) the less effective damage-control technology is.

3.2 Characterizing the worst-case pollution accumulation process

Eq. (20) specifies the worst-case misspecification of our model, given a value of $\theta > \sigma^2 b$. Substituting it into our robust pollution dynamics (7) yields

$$\mathrm{d}P(t) = \Big(m\bar{P} \underbrace{-\frac{\sigma^2}{\theta}\alpha_2(\theta,z)}_{\text{Effect 1}} + E - \Big[m + \underbrace{\frac{2\sigma^2}{\theta}\alpha_1(\theta,z)}_{\text{Effect 2}}\Big]P(t)\Big)\mathrm{d}t + \sigma\mathrm{d}B(t) \tag{28}$$

Eq. (28) points to two negative effects of model misspecification. First, there now exists an additional constant drift term (Effect 1) equal to

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

suggesting the presence of exogenous sources of pollution beyond those responsible for pre-industrial pollution stock $\bar{P}$. 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

$$\mathrm{d}P^*(t) = (m\bar{P} + E^* - mP^*(t) + \sigma \cdot v^*(t))\mathrm{d}t + \sigma\mathrm{d}B(t), \tag{29}$$

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

$$\mathrm{d}P^*(t) = -\Big[2\alpha_1(\theta,z)(\frac{1}{b} - \frac{\sigma^2}{\theta}) - m\Big]\Big(\frac{m\bar{P} + a + \alpha_2(\theta,z)(\frac{1}{b} - \frac{\sigma^2}{\theta})}{-[2\alpha_1(\theta,z)(\frac{1}{b} - \frac{\sigma^2}{\theta}) - m]} - P^*(t)\Big)\mathrm{d}t + \sigma\mathrm{d}B(t) \tag{30}$$

Stochastic differential equation (30) is an instance of the well-known Ornstein-Uhlenbeck process with parameters,

$$
\begin{aligned}
\mu(\theta,z) &= \frac{m\bar{P}+\frac{a}{b}+\alpha_2(\theta,z)(\frac{1}{b}-\frac{\sigma^2}{\theta})}{-[2\alpha_1(\theta,z)(\frac{1}{b}-\frac{\sigma^2}{\theta})-m]} = \bar{P}+\frac{a(m+\rho)}{b\Big(m^2+m\rho+gz(\frac{1}{b}-\frac{\sigma^2}{\theta})\Big)} \\
\xi(\theta,z) &= -\Big[2\alpha_1(\theta,z)(\frac{1}{b}-\frac{\sigma^2}{\theta})-m\Big] = \frac{\sqrt{(2m+\rho)^2+4gz(\frac{1}{b}-\frac{\sigma^2}{\theta})}-\rho}{2}.
\end{aligned}
\tag{31}
$$

As a result, we can establish the following:

Proposition 2 *Consider $\mu(\theta,z)$ and $\xi(\theta,z)$ as given by Eq. (31). Stochastic differential equation (30) 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

$$\mathbf{E}[P^*(\theta,z,t)] = \hat{P}_0 \mathrm{e}^{-\xi(\theta,z)t}+\mu(\theta,z)\Big[1-\mathrm{e}^{-\xi(\theta,z)t}\Big],$$

and variance

$$\mathbf{Var}[P^*(\theta,z,t)] = \frac{\sigma^2}{2\xi(\theta,z)}\Big[1-\mathrm{e}^{2\xi(\theta,z)t}\Big]^2.$$

(b) $\{P^*(\theta,z,t):t\geq 0\}$ has a stationary distribution that is $N\Big(\mu(\theta,z),\frac{\sigma^2}{2\xi(\theta,z)}\Big)$.

Proposition 2 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 $P^*(\theta,z,t)$, the entropy of our worst-case model misspecification has a closed-form expression:

$$
\begin{aligned}
R(Q^*(\theta,z)) &= \int_0^\infty \mathrm{e}^{-\rho t}\frac{1}{2}\mathbf{E}_{Q^*}[v^*(t)^2]\mathrm{d}t \\
&= \frac{\sigma^2}{2\theta^2}\int_0^\infty \mathrm{e}^{-\rho t}\Big[4\alpha_1^2(\theta,z)\Big((\mathbf{E}[P^*(\theta,z,t)])^2+\mathbf{Var}[P^*(\theta,z,t)]\Big) \\
&\quad +4\alpha_1(\theta,z)\alpha_2(\theta,z)\mathbf{E}[P^*(\theta,z,t)]+\alpha_2^2(\theta,z)\Big]\mathrm{d}t.
\end{aligned}
\tag{32}
$$

Thus, we are able to (via Proposition 1) directly associate an entropy bound $\eta^* = R(Q^*(\theta, z))$ to a given ambiguity parameter θ, such that the respective multiplier (10) and constraint (11) robust control problems admit identical solutions.

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 basic science suggest 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 mutiplier θ with technology adoption z, at initial pollution P_0, given by Eq. (21). Thus, at time 0, the policy maker wishes to solve the following optimization problem

$$\max_{z\in[0,1]} \qquad V(P_0, \theta, z) - \phi(z). \tag{33}$$

Lemma 2 *Suppose $P_0 \geq \bar{P}$ and consider optimization problem (33). There exists a unique optimal level of damage-control investment z, call it $z^*(\theta)$, that satisfies*

$$\begin{array}{ll}
(a) & \dfrac{\partial}{\partial z} V(P_0, \theta, z) > \phi'(z), \quad \text{for all } z \in [0, z^*(\theta)) \\
 & \dfrac{\partial}{\partial z} V(P_0, \theta, z^*(\theta)) = \phi'(z^*(\theta)), \\
 & \dfrac{\partial}{\partial z} V(P_0, \theta, z) < \phi'(z), \quad \text{for all } z \in (z^*(\theta), 1] \qquad \text{or,} \\
(b) & z^*(\theta) = 1
\end{array}$$

Proof. We distinguish between two cases.

Case 1.

$$\frac{\partial}{\partial z} V(P_0, \theta, 1) < \phi'(1). \tag{34}$$

Recall that ϕ is strictly decreasing and convex, and satisfies $\phi'(0) = -\infty$. This fact, in combination with Lemma 1, Inequality (26), and Inequality (34) establishes that $z^*(\theta)$ must satisfy (a).

Case 2.

$$\lim_{\theta\to\infty} \frac{\partial}{\partial z} V(P_0, \theta, 1) \geq \phi'(1). \tag{35}$$

In this case Lemma 1 and first-order conditions immediately imply $z^*(\theta) = 1$, in accordance with (b). ■

We are now ready to prove that optimal investment in damage-control technology is increasing in model uncertainty and thus consistent with the PP.

Theorem 1 *Suppose $P_0 \geq \bar{P}$. Optimal damage-control investment increases in model uncertainty. In other words, $z^*(\theta)$ is increasing in θ.*

Proof. Consider $\theta_2 > \theta_1$ and the associated optimal investment decisions $z^*(\theta_1)$ and $z^*(\theta_2)$. Suppose first that $z^*(\theta_1) < 1$. Then Lemma 2 implies that $z^*(\theta_1)$ uniquely satisfies

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

Lemma 1 further implies that

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

Consequently, Lemma 2 leads to the following inequality:

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

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

Suppose now that $z^*(\theta_1) = 1$ so that taking derivatives we obtain

$$\frac{\partial}{\partial z} V(P_0, \theta_1, 1) \geq \phi'(z^*(\theta_1)).$$

By similar reasoning we can establish $\frac{\partial}{\partial z} V(P_0, \theta_2, 1) \geq \phi'(1)$ implying $z^*(\theta_2) = 1$. ■

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

Theorem 2 *Suppose $P \geq \bar{P}$ and consider a neighborhood of θ, say $(\theta_{\min}, \theta_{\max}] \subseteq (b\sigma^2, \infty]$. If $z^*(\theta)$ satisfies*

(a)

$$\frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta) > \frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta)) - \frac{2\beta_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in (\theta_{\min}, \theta_{\max}], \tag{36}$$

then robustly-optimal emissions $E^*(P)$ are unambiguously decreasing in θ in $(\theta_{\min}, \theta_{\max}]$;

(b)

$$\frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta) < \frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta))}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in (\theta_{\min}, \theta_{\max}], \tag{37}$$

then robustly-optimal emissions $E^*(P)$ are unambiguously increasing in θ in $(\theta_{\min}, \theta_{\max}]$;

(c)

$$\frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta))}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))} < \frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta) < \frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta)) - \frac{2\alpha_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))}, \quad \theta \in (\theta_{\min}, \theta_{\max}] \tag{38}$$

then robustly-optimal emissions $E^*(P)$ will be decreasing in θ for $\theta \in (\theta_{\min}, \theta_{\max}]$ if and only if current pollution levels are high enough.

Proof. Consider θ and the associated optimal $z^*(\theta)$. We begin by showing that the optimal solution of optimization problem (33) is such that the values of $\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_1(\theta, z^*(\theta))$ and $\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_2(\theta, z^*(\theta))$ can be positive or negative. In particular, we prove the following:

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_1(\theta, z^*(\theta)) < 0 \quad \Leftrightarrow \quad \frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta) > \frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta))}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))} \tag{39}$$

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_2(\theta, z^*(\theta)) < 0 \quad \Leftrightarrow \quad \frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta) > \frac{-\frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta)) - \frac{2\beta_1^2(\theta, z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))}. \tag{40}$$

We begin with (39) and consider $\beta_1(\theta, z^*(\theta))$. The result immediately follows from differentiating with respect to θ and recalling the negative sign of $\frac{\partial \beta_1}{\partial z}(\theta, z)$:

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_1(\theta, z^*(\theta)) = \frac{\partial \beta_1}{\partial \theta}(\theta, z^*(\theta)) + \frac{\partial \beta_1}{\partial z}(\theta, z^*(\theta))\frac{\mathrm{d}z^*}{\mathrm{d}\theta}(\theta). \tag{41}$$

Moving on to (40), we refer to Eq. (22). Straightforward differentiation establishes that

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_2(\theta, z^*(\theta)) = \frac{2(\frac{a}{b} + m\bar{P})\Big(\frac{\mathrm{d}}{\mathrm{d}\theta}\beta_1(\theta, z^*(\theta))(\rho + m) + \frac{2\beta_1^2(\theta,z^*(\theta))\sigma^2}{\theta^2}\Big)}{\Big(2\beta_1(\theta, z^*(\theta))(\frac{\sigma^2}{\theta} - 1) + \rho + m\Big)^2} \tag{42}$$

The result now may be arrived at through Eqs. (41) and (42).

The theorem now follows immediately from Expressions (39) and (40), and the fact that (as Eq. (27) suggests) $E^*(\theta, z, P) = \frac{1}{b}[a + 2\beta_1(\theta, z)(P - \bar{P}) + \beta_2(\theta, z)]$. ■

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

$$\frac{\partial\beta_1}{\partial\theta}(\theta, z) \geq 0, \quad \frac{\partial\beta_1}{\partial\theta}(\theta, 0) = 0, \tag{43}$$

while Lemma 1 implies that

$$\frac{\partial\beta_1}{\partial z}(\theta, z) < 0, \quad \frac{\partial^2\beta_1}{\partial\theta\partial z}(\theta, z) > 0, \quad \frac{\partial\beta_1}{\partial z}(\theta, 0) = -\frac{g}{4m + 2\rho}. \tag{44}$$

Therefore, the conditions of Theorem 2 are not generically false so that it is, theoretically, possible for emissions to increase as uncertainty goes up. Moreover, Eqs. (22), (43), and (44) imply that the right-hand-side of Eq. (36) is increasing in $z^*(\theta)$ and satisfies

$$\lim_{z^*(\theta)\to 0} \frac{-\frac{\partial\beta_1}{\partial\theta}(\theta, z^*(\theta)) - \frac{2\beta_1^2(\theta,z^*(\theta))\sigma^2}{\theta^2(\rho+m)}}{\frac{\partial\beta_1}{\partial z}(\theta, z^*(\theta))} = 0.$$

Hence, we arrive at the following corollary of Theorem 2.

Corollary 1 *The right-hand side of Eq. (36) is increasing in $z^*(\theta)$, and vanishes at $z^*(\theta) = 0$. Hence, fow 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 subtantial and 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*.

5 Numerical Results

5.1 Preliminaries

In this section we perform a numerical exercise that provides some context for the theoretical results. We focus on the following family of cost functions that is consistent with our model assumptions

$$\phi(z;k) = k\left(\frac{1}{z^2} - 1\right), \quad k > 0, \tag{45}$$

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 *Suppose $P_0 \geq \bar{P}$. Fix a level of uncertainty θ and consider a family of optimization problems (33), parametrized according to Eq. (45).*

(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) Suppose $P \geq \bar{P}$. Optimal emissions $E^*(P;k)$ 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 at the least. 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.

5.2 An application to climate change economics

To make the analysis concrete, we focus on a climate-change application of our model and calibrate the relevant parameters according to Karp and Zhang [20]. The standard deviation of the carbon accumulation process, σ, is calibrated on data compiled by the US Dept of Commerce's National Oceanic and Atmospheric Administration (NOAA).[11] Table 1 summarizes the values of all model parameters.

Damage Control. We already know from Theorem 1 that optimal damage-control investment is increasing in uncertainty, i.e., that $z^*(\theta;k)$ is increasing in θ, for all cost functions (45). Indeed this can be readily seen in Figure 1, in which optimal damage-control investment is plotted as a function of θ for a variety of cost functions. Figure 1 further illustrates Proposition 3: given a level of uncertainty θ, optimal damage-control investment is decreasing in the cost of technology. The chosen values of k lead to a wide spectrum of damage-control investments, ranging from the very aggressive to absolutely zero. On the one extreme, when $k = 1000$ and damage-control technology is very cheap, investment is very high so that around 90% of damages are directly reduced. As k increases, this investment becomes smaller and smaller, until we reach $k = 25000$, at which point there will be positive investment in damage

[11] See NOAA's website on Trends in Carbon Dioxide at http://www.esrl.noaa.gov/gmd/ccgg/trends/global.html. Our value of σ is derived from Mauna Loa data on annual mean growth rates of CO_2 for the period 1959 to 2010, which can be found at: ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_gr_mlo.txt.

Parameter	Description	Value	Unit
P_0	base year pollution stock	781	GtC
$\bar{P}$	pre-industrial pollution stock	590	GtC
g	slope of marginal damage	0.0223	10^9\$ / (GtC)^2
a	intercept of marginal benefit	224.26	\$ / tC
b	slope of marginal benefit	1.9212	10^9\$ / (GtC)^2
m	carbon decay rate	0.0083	scalar
σ	carbon standard deviation	0.2343	GtC
ρ	pure rate of time preference	.03	scalar

Table 1: Calibration of model parameters based on Karp and Zhang [20] and NOAA (see text). When there is no uncertainty or damage control investment (i.e., when $\theta = \infty$ and $z = 1$), the calibration results in a steady-state carbon stock of approximately 965 GtC (453 ppm CO_2) that, according to prevailing climate science, is more or less consistent with a 2^oC warming stabilization target.

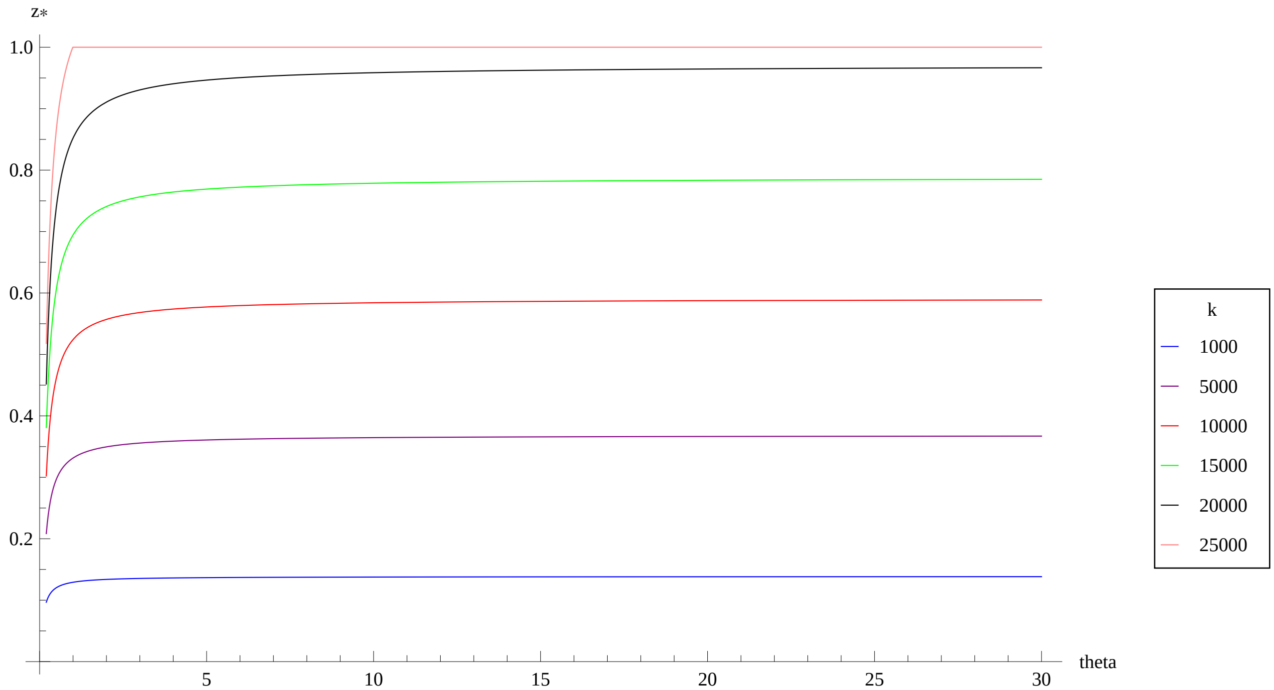

Figure 1: Optimal damage-control investment as a function of θ for different k.

control only for extreme, and physically implausible, levels of misspecification.[12] A consistent trend for all k is that $\frac{\mathrm{d}z^*(\theta;k)}{\mathrm{d}\theta}$ is decreasing in θ, with very large values close to the origin that fast taper off towards 0 for $\theta > 1$. This result suggests that the magnitude of model misspecification is important primarily when uncertainty is high; when this is not the case, optimal investment in damage control technology is not all that sensitive to the degree of model misspecification.

Mitigation. While our choice of k does not affect the PP with regard to damage-control investment, this is not true in the case of mitigation. Instead we observe an ambiguous relationship, as predicted by Theorem 2. Throughout the following exercises we calibrate the multiplier θ, and by extension the relative entropy bound η, by carefully considering the worst-

[12]Figure 1 shows that for $k = 25000$ we have $z^*(\theta) < 1$ only for $\theta \leq 1$. The worst-case model misspecification corresponding to such low values of θ implies a (nonsensical) *negative* carbon decay rate of less than -0.00128.

case misspecified dynamics our choices lead to. In particular, we focus on Eq. (28)'s Effects 1 (an increase in exogenous sources of carbon) and 2 (a decrease in the natural decay rate of carbon) and pick values of θ that provide reasonable bounds on their worst-case *percentage* deviations from the benchmark case. For, recall that (by Propositions 1 and 2) when we solve the multiplier problem for a particular choice of θ, this is akin to finding a robust policy for all *all* probability models having relative entropy less than the distorted model in which, concentrating on Effects 1 and 2 and Eq. (28), we observe percentage deviations of

$$\Bigg[\underbrace{\frac{\text{Effect } 1(\theta, z^*(\theta))}{m\bar{P}}}_{\text{\% increase of exogenous pollution}}, \quad \underbrace{\frac{\text{-Effect } 2(\theta, z^*(\theta))}{m}}_{\text{\% decrease of carbon decay rate}} \Bigg]$$

from the benchmark (4). [Note that this entropy will equal $R(Q^*(\theta, z^*(\theta)))$ as given by Eq. (32).]

Bearing the above in mind, we set θ in such a way as to provide sensible values for the following expression:

$$\begin{aligned} \text{Deviation}(\theta) \quad &\equiv \quad \%\text{Eff1}(\theta, z^*(\theta)) + \%\text{Eff2}(\theta, z^*(\theta)) = \frac{\text{Effect } 1(\theta, z^*(\theta))}{m\bar{P}} + \frac{\text{-Effect } 2(\theta, z^*(\theta))}{m} \\ &= \quad \frac{-\frac{\sigma^2}{\theta}\alpha_2(\theta, z^*(\theta))}{m\bar{P}} + \frac{-\frac{2\sigma^2}{\theta}\alpha_1(\theta, z^*(\theta))}{m}. \end{aligned} \tag{46}$$

Eq. (46) grounds our choice of θ to the underlying physics of carbon accumulation through an aggregation of Effects 1 and 2, and allows for a systematic comparison of model results across different cost functions. In what follows we choose values of θ that, using the formula given by Eq. (46), lead to a Deviation(θ) of 0%, 10%, 50%, 100%, and 200%.

We focus on the two lowest cost functions that were presented in Figure 1, corresponding to $k = 1000$ and $k = 5000$. Figures 2 and 3 illustrate part (c) of Theorem 2 and show how, in both cases, a reversal of the PP with regard to mitigation is in principle possible for high enough levels of current carbon stock P. The substantive difference between the two lies in the *probability* of this reversal ever being observed. When $k = 1000$ we claim that this probability is high, whereas for $k = 5000$ it is negligible.

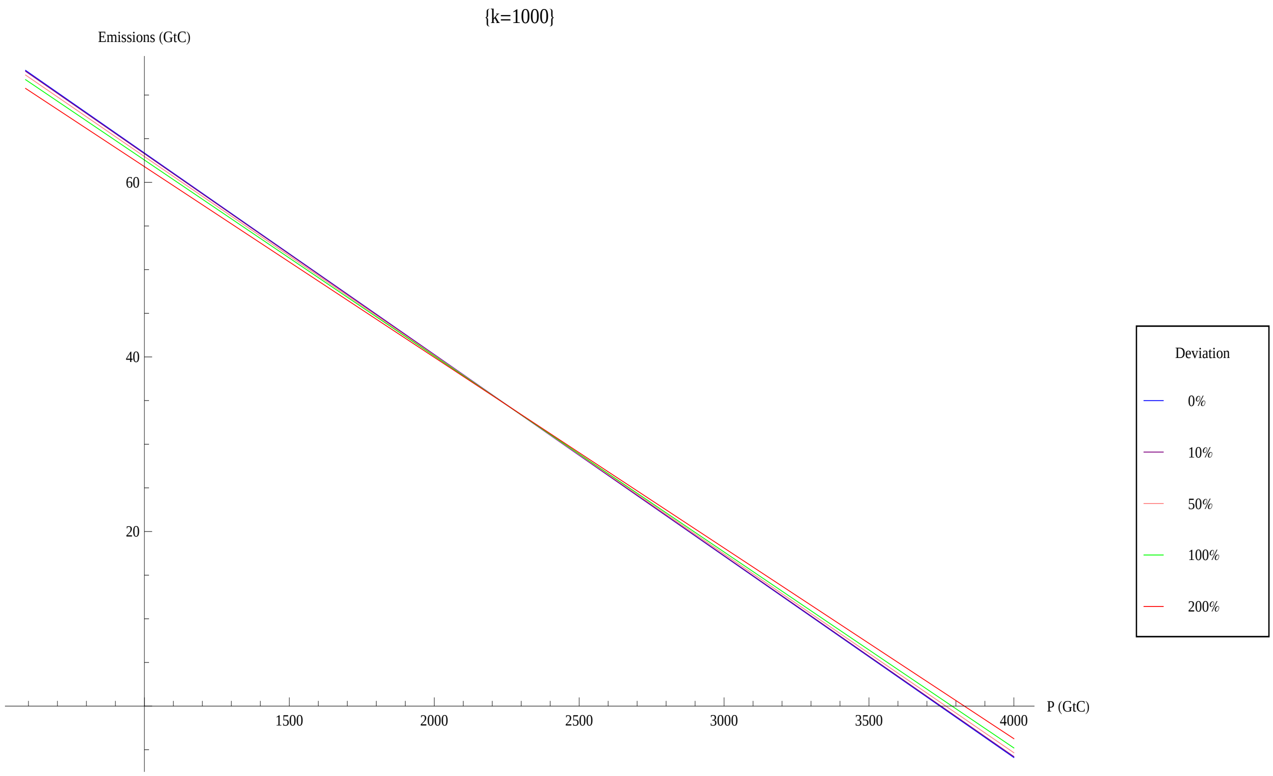

Figure 2: Robust emissions policy for $k = 1000$ and different levels of model misspecification.

θ	Deviation	%Eff 1	% Eff 2	η^*	z^*	$\mathbf{E}[P^*]$	$\mathbf{E}[P^R_{\min}]$
10^6	0	0	0	0	.1386	2910.5	2910.5
9.4	.1	.069	.031	216	.1376	2947	2860
1.91	.5	.350	.150	5431	.1336	3097.5	2720.9
.96	1	.712	.297	22525	.1288	3301.9	2452.9
.493	2	1.445	.563	93426	.1200	3754	2070.7

Table 2: $k = 1000$. $P^R_{\min}$ denotes approximately minimal steady-state carbon stock under the robust policy.

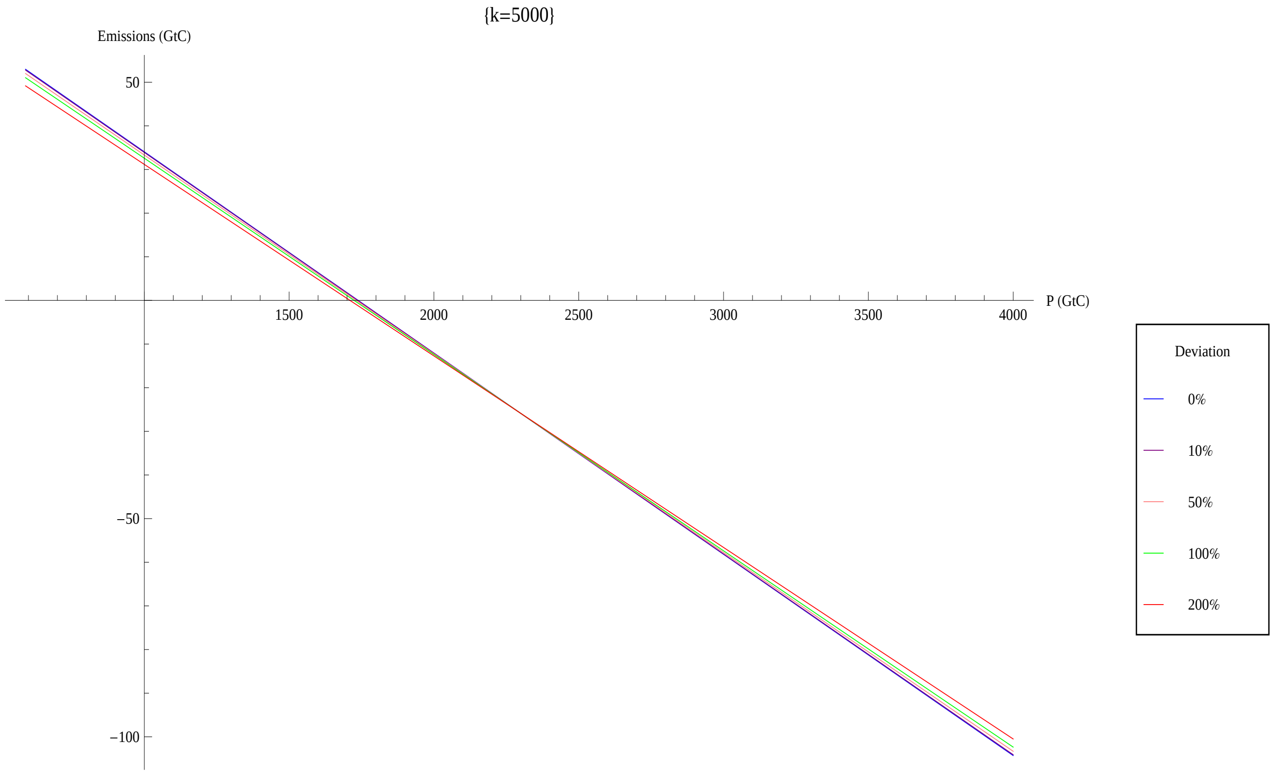

Figure 3: Robust emissions policy for $k = 5000$ and different levels of model misspecification.

θ	η^*	$\mathbf{E}[P^R_{\max}]$	$\mathbf{E}[P^{NR}_{\max}]$	$\mathbf{E}[P^R_{\infty}]$	$\mathbf{E}[P^{NR}_{\infty}]$
9.4	216	2965.5	2965.5	2911.8	2910.5
1.91	5431	3188.6	3180.8	2917	2910.5
.96	22525	3455	3429.6	2923.7	2910.5
.493	93426	3987.3	3896.2	2937.3	2910.5

Table 3: $k = 1000$. $P^{R(NR)}_{\max}$ denote approximately maximal steady-state carbon under robust (non-robust) policies. P^R_{∞} denote steady-state carbon levels of robust policies under model certainty.

θ	Deviation	% Eff 1	%Eff 2	η^*	z^*	$\mathbf{E}[P^*]$
10^6	0	0	0	0	.3683	1563.3
13.7	.1	.058	.042	169	.3656	1577.2
2.785	.5	.350	.150	4250	.3548	1634.5
1.41	1	.712	.297	17395	.3419	1710.7
.727	2	1.445	.563	71794	.3181	1876.9

Table 4: $k = 5000$.

So why is the PP violated with high probability in the case of $k = 1000$? Notice from Figure 2 that optimal emissions are increasing in model misspecification in the range of about $P \geq 2300$. Now, take a look at column 7 of Table 2 depicting steady-state values for the expected value of worst-case carbon stock levels. When θ is very high and there model certainty, the expected value of the carbon stock is 2910.5 GtC, much higher than 2300. For positive levels of model misspecification expected carbon levels corresponding to the worst-case model misspecification pollution stock are higher than the benchmark 2910.5 GtC and significantly higher than 2300.[13]

But these values correspond to worst-case outcomes and so do not necessarily provide adequate insight into the probability of exceeding 2300 GtC, given the set of probability models (as defined by relative entropy bounds of Eq. 9) we are seeking robustness over. So for each of our misspecifications, we compute a plausible approximation of the *lowest* possible expected level of steady-state pollution given the relevant relative entropy bounds η^*. We do this in the following way. Given our choices of θ we first consider the optimal damage-control investment $z^*(\theta)$ and the robust feedback policy $E^*(\theta, z^*(\theta))$ and entropy bound $\eta^*(\theta, z^*(\theta)) \equiv \eta^*$ they lead to (given by Eqs. (19) and (32), respectively). Consequently,

[13]Steady-state variance levels are very low compared to mean values (generally less than 1GtC) and therefore, in light of Proposition 2, unimportant from a practical standpoint. This remains true for all results reported in Tables 2, 3 and 4.

we solve for the model misspecification (in the notation of Sections 2 and 3, the control variable v) that approximately minimizes expected steady-state carbon levels, subject to the feedback rule E^* and entropy constraint η^*. Given the broad range of possible model misspecifications and the generic intractability of stochastic differential equations, performing this calculation is not in principle a simple matter. But fortunately our problem structure justifies concentration on a specific class of tractable model missprecifications so that the resulting optimization problem can be solved efficiently.[14] The outcome of these computations appears in column 8 of Table 2. We see that for Deviations of 10%, 50%, and 100% (i.e., $\theta \in \{9.4, 1.91, .94\}$) even approximately minimal expected pollution levels will be significantly higher than 2300 GtC in steady-state. Moreover, when Deviation(θ) is equal to 200% they will be around 2070 GtC, only modestly below the threshold.

What this all implies is that, for all the chosen values of θ, it is very likely that *en route to* a steady state, carbon levels will exceed 2300 GtC. Hence, we will, with substantial probability, find ourselves in a range of P for which we observe a reversal of the PP with respect to mitigation.

Further indications that robust policies are not necessarily precautionary when $k = 1000$ can be seen by comparing them to their non-robust counterpart. First, using a similar approach as the one employed for the aforementioned minimizations (described in section 2 of the Appendix) we compute approximate values for the *highest* possible steady-state pollution levels subject to the relevant entropy constraints, under both the robust policies (obtained by plugging in appropriate values of θ and $z^*(\theta)$ into Eq. 19) and the non-robust policy (obtained by plugging in $\theta = \infty$ and $z^*(\infty)$ into Eq. 19). These results appear in columns 3 and 4 of Table 3 and demonstrate that, given the relevant entropy bounds, robust policies consistently lead to *higher* worst-case expected pollution. A second sign of the non-precautionary character of robust policies can be seen when fears of model misspecification are

[14]Proposition 2 is especially helpful in performing these calculations. For details the reader is referred to the section 2 of the Appendix. All computations were performed in Mathematica.

unfounded and we have model certainty. In particular, we calculate the expected steady-state carbon stock levels that the robust policies $E^*(\theta, z^*(\theta))$ lead to when, unbeknownst to the policy maker, $\theta = \infty$ and the benchmark model (4) uniquely captures carbon dynamics. The results of our computations appear in columns 5 of Table 3. Again, robust policies consistently lead to higher steady-state carbon stock compared to their non-robust counterparts, and this difference is increasing in the perceived (yet imaginary) degree of uncertainty.

When $k = 5000$ the situation is markedly different. Notice from Figure 3 that emissions again begin being increasing in model uncertainty around $P = 2300$. Now, take a look at column 7 of Table 4 depicting steady-state values for the expected value of worst-case carbon-stock levels. Even when model misspecification is at its highest level, corresponding to a 200% joint miscalculation of Effects 1 and 2, they will not exceed 1880 GtC. Moreover, a similar computation as the one that was done for $k = 1000$ and reported in column 3 of Table 3 establishes that, within the relevant relative maximal entropy bound of 71794, expected steady state carbon stocks peaks at around 1894 GtC. Thus, for all our chosen values of θ, it is highly unlikely that carbon levels will ever exceed the threshold of 2300 GtC. We conclude that when $k = 5000$, even though theoretically possible, the probability of ever observing a reversal of the PP with respect to mitigation is negligible.

Results for $k \in \{10000, 15000, 20000\}$ are qualitatively similar to those for $k = 5000$ and omitted for brevity.[15] When $k = 25000$, the cost of damage-control is so high, and therefore investment in it so low and/or non-existent, that the reversal of the PP is not even in principle possible. Indeed, as soon as z^* hits 1 and the derivative $\frac{\mathrm{d}z^*}{\mathrm{d}\theta}$ equals 0 (see Figure 1), common sense, (as well as, more formally, Theorem 2) suggest that it becomes a mathematical impossibility.

We end this section by generally noting that, as Figures 2 and 3 show, robust policies do not seem to be very sensitive to changes in θ. This effect is a function of our model parameters (e.g. the low value of σ) as well as the high investments in damage control,

[15] Graphs available upon request.

which, as Eqs. (16) and (17) suggest, serves to temper differences in θ.

6 Conclusion

The present paper analyzed optimal pollution control policy under Knightian uncertainty by adopting the robust control framework of Hansen and Sargent [12]. Allowing for a one-time investment in 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 robust emissions policies can be increasing in model uncertainty.

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 precautious now by undertaking large damage-control investment, and not be particularly precautious with respect to future mitigation policy. 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. The theoretical results are consequently applied to a linear-quadratic model of climate change, calibrated by Karp and Zhang [20]. In our simulations we take pains to carefully calibrate the uncertainty parameter of our model and provide a conceptual link to the actual dynamic process of carbon accumulation. The methods we employ build on the preceding theoretical analysis and may be, at least in our view, of independent interest for robust control applications. Our main policy-relevant finding is that emissions can be increasing in uncertainty only when damage-control technology is extremely and most probably unrealistically cheap. Thus, at least within the context of this numerical model, we do not expect our more "controversial" theoretical findings to be of much practical relevance.

This work suggests several interesting avenues for future research. A more complete treatment of the issues presented here would extend the basic model to incorporate dynamic damage-control investment, more intricate pollution dynamics, and lower bounds on emissions that would reflect concerns about irreversibility.

Appendix

1. Monotonicity properties of $\frac{\partial \beta_3}{\partial z}$

Recall that

$$\Delta(\theta, z) = \sqrt{(2m+\rho)^2 + 4gz(\frac{1}{b} - \frac{\sigma^2}{\theta})}$$

and that we want to show that $\frac{\partial \beta_3}{\partial z}$ is increasing in both θ and z. After simplifying we obtain:[16]

$$\begin{aligned}
\frac{\partial^2 \beta_3}{\partial z \partial \theta}(\theta, z) &= \frac{f_1(\theta, z)}{f_2(\theta, z)}, \text{ where} \\
f_1(\theta, z) &= 8g^2\sigma^2 z(\theta - b\sigma^2)\Big[2g\rho\theta z\sigma^2 b\Delta(\theta, z) + 2g^2\sigma^2 z^2(\theta - b\sigma^2) + m\rho\theta g\sigma^2 zb + 4gp^2\sigma^2\theta zb \\
&+ 4m^2\theta g\sigma^2 z\Big] + 8g^2\sigma^2 z\Big[2a^2(m+\rho)^2\theta^2(\rho + 4\Delta(\theta, z)) + b^2\sigma^2\Big[2m^4\theta^2 + 4m^3\rho\theta^2 + \rho^4\theta^2 + \rho^3\theta^2\Delta(\theta, z) \\
&+ 2m\rho\theta\big[2\rho^2\theta + \rho\theta\Delta(\theta, z)\big] + 2m^2\theta\big[3\rho^2\theta + \rho\theta\Delta(\theta, z)\big]\Big]\Big], \\
f_2(\theta, z) &= b\rho\theta^3\Delta(\theta, z)\big[\rho + \Delta(\theta, z)\big]^4\big[4gz(\theta - b\sigma^2) + 4m^2\theta b + 4m\rho b\theta + b\rho^2\theta\big]
\end{aligned}$$

$$\begin{aligned}
\frac{\partial^2 \beta_3}{\partial z^2}(\theta, z) &= \frac{g_1(\theta, z)}{g_2(\theta, z)}, \text{ where} \\
g_1(\theta, z) &= \frac{\theta - b\sigma^2}{z} f_1(\theta, z) \\
g_2(\theta, z) &= \frac{\rho}{\theta} f_2(\theta, z)
\end{aligned}$$

Since θ satisfies $\theta > \sigma^2 b$, it is clear that all of the above are strictly postive. ■

2. Minimization & maximization of steady-state pollution levels given entropy constraints

Given $\theta > \sigma^2 v$, consider the optimal damage-control decision $z^*(\theta)$ and the consequent optimal emissions feedback policy $E^*(P) = \frac{1}{b}\Big[a + \alpha_2(\theta, z^*(\theta)) + 2\alpha_1(\theta, z^*(\theta))P\Big]$. These will lead to a relative entropy bound $\eta^*(\theta, z^*(\theta)) \equiv \eta^*$ given by Eq. (32). The optimization problem we ideally wish to solve is the following

$$\min_{v} \quad \mathbf{E}\big[\lim_{t\to\infty} P(t)\big]$$

[16] Mathematica output available upon request.

$$
\begin{aligned}
\text{subject to:} \quad & \mathrm{d}P(t) = \big[E^*(P(t)) - m(P(t)-\bar{P}) + \sigma v\big]\mathrm{d}t + \sigma \mathrm{d}B(t) \\
& \int_0^\infty \frac{1}{2}\mathrm{e}^{-\rho t}\mathbf{E}[v(t)^2]\mathrm{d}t \le \eta^*, \quad P(0) = P_0.
\end{aligned} \tag{47}
$$

We conjecture that there exists an, at least approximately, optimal solution to the optimization problem (47) that is linear in P so that $v^*(P) = \frac{1}{\sigma}\big(\gamma_1 - \gamma_2 P\big)$ for some γ_1 and γ_2. While a formal investigation of this statement is beyond the scope of the current paper, we base our intuition on the fact that for any convex and quadratic function $f(\cdot)$, linear feedback policy $E(P)$, and discount rate $\tilde{\rho}$, the following optimization problem:

$$
\begin{aligned}
\min_{v} \quad & \mathbf{E}\int_0^\infty \mathrm{e}^{-\tilde{\rho}t} f(P(t))\mathrm{d}t \\
\text{subject to:} \quad & \mathrm{d}P(t) = \big[E(P(t)) - m(P(t)-\bar{P}) + \sigma v\big]\mathrm{d}t + \sigma \mathrm{d}B(t) \\
& \frac{1}{2}\int_0^\infty \mathrm{e}^{-\tilde{\rho}t}\mathbf{E}\big[v(t)^2\big]\mathrm{d}t \le \eta^*, \quad P(0) = P_0,
\end{aligned} \tag{48}
$$

has an optimal solution v that is linear in the state variable.[17]

With the above in mind, we return to problem (47). To make our domain at least somewhat realistic we restrict ourselves to misspecifications of the type

$$
\begin{aligned}
& v(P) = \frac{1}{\sigma}\left[\gamma_1 - \gamma_2 P\right], \quad \text{where} \\
& -m\bar{P} \le \gamma_1 \le 10m\bar{P}, \quad -m \le \gamma_2 \le 5m.
\end{aligned} \tag{49}
$$

Plugging in the robust feedback policy $E^*(P)$ and the above choice (49) for $v(\cdot)$ to the stochastic differential equation (7), once again leads to an Ornstein-Uhlenbeck process $\{P(t;\gamma_1,\gamma_2) : t \ge 0\}$ with parameters:

$$
\begin{aligned}
\mu(\theta,\gamma_1,\gamma_2) &= \frac{m\bar{P} + \frac{a+\alpha_2(\theta,z^*(\theta))}{b} + \gamma_1}{-\frac{2\alpha_1(\theta,z^*(\theta))}{b} + m + \gamma_2} \\
\xi(\theta,\gamma_1,\gamma_2)) &= -\frac{2\alpha_1(\theta,z^*(\theta))}{b} + m + \gamma_2.
\end{aligned}
$$

This leads to a steady-state distribution that is again $N\bigg(\mu(\theta,\gamma_1,\gamma_2), \frac{\sigma^2}{2\xi(\theta,\gamma_1,\gamma_2))}\bigg)$.[18] The values of $\mathbf{E}[P^R_{\min}]$ quoted in Table 2 correspond to the optimal values of the following minimization problem:

$$
\min_{\gamma_1,\gamma_2} \quad \mu(\theta,\gamma_1,\gamma_2)
$$

[17]To see this result, note that optimization problem (48) is again linear-quadratic so that the reasoning of Sections 2 and 3 applies. The only difference is that here we only minimize over v without maximizing over E.

[18]Note that setting $\gamma_1 = -\frac{\sigma^2}{\theta}\alpha_2(\theta,z^*(\theta))$ and $\gamma_2 = \frac{2\sigma^2}{\theta}\alpha_1(\theta,z)$ recovers our worst-case model misspecification as per Eq. (28).

$$\text{subject to:} \quad \int_0^\infty \frac{1}{2\sigma^2} \mathrm{e}^{-\rho t} \mathbf{E}\big[\gamma_1 + \gamma_2 P(t;\gamma_1,\gamma_2)\big]^2 \mathrm{d}t \le \eta^*(\theta)$$

$$-m\bar{P} \le \gamma_1 \le 10m\bar{P}, \quad -m \le \gamma_2 \le 5m, \quad P(0) = P_0 \tag{50}$$

Equivalent reasoning applies for the maximization problem and the values of $\mathbf{E}[P^R_{\max}]$ quoted in Table 3. ■

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] M. Funke and M. Paetz (2010), "Environmental Policy under Model Uncertainty: A Robust Optimal Control Approach," *Climatic Change*, 107, 225–239.

[8] I. Gilboa and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

[9] C. Gollier, J. Gierlinger (2008), "Socially efficient discounting under ambiguity aversion," Working paper.

[10] F. Gonzalez (2008), "Precautionary Principle and Robustness for a Stock Pollutant with Multiplicative Risk," *Environmental and Resource Economics*, 41, 25–46.

[11] L.P. Hansen, T. Sargent (1995), "Discounted linear exponential quadratic Gaussian control", *IEEE Transactions on Automatic Control,* 40, 968971.

[12] L. P. Hansen and T. Sargent (2001), "Robust Control and Model Uncertainty," *American Economic Review P & P*, 91, 60-66.

[13] L. P. Hansen and T. Sargent (2007), "Recursive Robust Estimation and Control Without Commitment," *Journal of Economic Theory* 136, 1–27.

[14] L. P. Hansen and T. Sargent (2008), *Robustness* Pinceton University Press.

[15] L. P. Hansen, T. Sargent, G. Turhumambetova, and N. Williams (2006), "Robust Control and Model Misspecification," *Journal of Economic Theory*, 128, 45-90.

[16] M. Hoel and L. Karp (2001), "Taxes Versus Quotas for a Stock Pollutant under Multiplicative Uncertainty," *Journal of Public Economics*, 82, 91–114.

[17] D. Jacobson (1973), "Optimal Stochastic Linear Systems with Exponential Performance Criteria and their Relation to Deterministic Differential Games," *IEEE Transactions on Automatic Control*, AC-18, 124–131.

[18] M. R. James (1992), "Asymptotic Analysis of Nonlinear Stochastic Risk Sensitive Control and Differential Games," *Mathematics of Control, Signals and Systems*, 5, 401–417.

[19] M. R. James and R. Elliott (1992), "Risk-Sensitive Control and Dynamic Games for Partially Observed Discrete-Time Nonlinear Systems," *IEEE Transactions on Automatic Control,* 39, 780–792.

[20] L. Karp and J. Zhang (2006), "Regulation with Anticipated Learning about Environmental Variables," *Journal of Environmental Economics and Management*, 51, 259-279.

[21] P. Klibanoff, M. Marinacci, and S. Mukerji (2005), "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73, 1849–1892.

[22] P. Klibanoff, M. Marinacci, and S. Mukerji (2009), "Recursive Smooth Ambiguity Preferences," *Journal of Economic Theory*, 144, 930–976.

[23] F. Knight (1921), *Risk, Uncertainty, and Profit* Houghton Mifflin, USA.

[24] D. Lemoine and C. Traeger, "Tipping Points and Ambiguity in the Integrated Assessment of Climate Change," Working paper.

[25] A. Millner, S. Dietz, and G. Heal (2010) "Ambiguity and climate policy," Working paper.

[26] J. von Neumann and O. Morgenstern. (1944), *Theory of Games and Economic Behavior*, Princeton University Press.

[27] W. Nordhaus (2008), *A Question of Balance: Economic Modeling of Global Warming*, Yale University Press.

[28] C. Roseta-Palma and A. Xepapadeas (2004) “Robust control in water management,” *Journal of Risk and Uncertainty* 29, 21–34.

[29] L. J. Savage (1954), *The Foundations of Statistics*, Wiley, New York.

[30] N. Stern (2007), *Stern Review: The Economics of Climate Change*, Cambridge University Press

[31] 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.

[32] N. Treich (2010), “The Value of a Statistical Life under Ambiguity Aversion,” *Journal of Environmental Economics and Management*, 59, 15-26.

[33] G. Vardas and A. Xepapadeas (2010), “Model Uncertainty, Ambiguity, and the Precautionary Principle: Implications for Biodiversity Management,” *Environmental and Resource Economics*, 45, 379–404.

[34] M. Weitzman (2009),“On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*, 91, 1–19

[35] M. Weitzman (2010), “What is the “damages function” for global warming -and what difference might it make?” *Climate Change Economics* 1, 57-69.

[36] P. Whittle (1981), “Risk sensitive linear-quadratic control problems,” *Advances in Applied Probability* 13, 764–777.

[37] A. Xepapadeas (2011), “The Cost of Ambiguity and Robustness in International Pollution Control,” Working paper.