

11-13-2016


Working Paper Number 16011

Optimal Climate Policy When Damages are Unknown

Ivan Rudik

Iowa State University, irudik@iastate.edu

Follow this and additional works at: http://lib.dr.iastate.edu/econ_workingpapers

 Part of the [Agricultural and Resource Economics Commons](#), [Climate Commons](#), [Natural Resources Management and Policy Commons](#), and the [Public Economics Commons](#)

Recommended Citation

Rudik, Ivan, "Optimal Climate Policy When Damages are Unknown" (2016). *Economics Working Papers*. 16011.
http://lib.dr.iastate.edu/econ_workingpapers/16

Iowa State University does not discriminate on the basis of race, color, age, ethnicity, religion, national origin, pregnancy, sexual orientation, gender identity, genetic information, sex, marital status, disability, or status as a U.S. veteran. Inquiries regarding non-discrimination policies may be directed to Office of Equal Opportunity, 3350 Beardshear Hall, 515 Morrill Road, Ames, Iowa 50011, Tel. 515 294-7612, Hotline: 515-294-1222, email eooffice@mail.iastate.edu.

This Working Paper is brought to you for free and open access by the Iowa State University Digital Repository. For more information, please visit lib.dr.iastate.edu.

Optimal Climate Policy When Damages are Unknown

Abstract

Integrated assessment models (IAMs) are economists' primary tool for analyzing the optimal carbon tax. Damage functions, which link temperature to economic impacts, have come under fire because of their assumptions that may produce significant, and ex-ante unknowable misspecifications. Here I develop novel recursive IAM frameworks to model damage uncertainty. I decompose the optimal carbon tax into channels capturing parametric damage uncertainty, learning, and misspecification concerns. Damage learning and using robust control to guard against potential misspecifications can both improve ex-post welfare if the IAM's damage function is misspecified. However, these ex-post welfare gains may take decades or centuries to arrive.

Disciplines

Agricultural and Resource Economics | Climate | Natural Resources Management and Policy | Public Economics

Optimal Climate Policy When Damages are Unknown

Ivan Rudik*

Department of Economics and Center for Agricultural and Rural Development
Iowa State University
479 Heady Hall, Iowa State University, Ames, IA, 50011-1070, USA
irudik@iastate.edu

November 13, 2016

Integrated assessment models (IAMs) are economists' primary tool for analyzing the optimal carbon tax. Damage functions, which link temperature to economic impacts, have come under fire because of their assumptions that may produce significant, and ex-ante unknowable misspecifications. Here I develop novel recursive IAM frameworks to model damage uncertainty. I decompose the optimal carbon tax into channels capturing parametric damage uncertainty, learning, and misspecification concerns. Damage learning and using robust control to guard against potential misspecifications can both improve ex-post welfare if the IAM's damage function is misspecified. However, these ex-post welfare gains may take decades or centuries to arrive.

JEL: H23, Q54, Q58

Keywords: climate, integrated assessment, damages, deep uncertainty, robust control, sparse grid, learning, social cost of carbon, carbon tax

*This paper was originally circulated as "Targets, Taxes, and Learning: Optimizing Climate Policy Under Knightian Damages." I am immensely grateful for advice and support from Derek Lemoine. I also benefited from comments and discussions with Dave Kelly, Lint Barrage, Stan Reynolds, Alex Hollingsworth, Price Fishback, and Quinn Weninger as well as participants at the Western Economics Association Meeting, the SCRiM Summer School, the CU Environmental and Resource Economics Workshop, The Heartland Environmental and Resource Economics Workshop, The AERE Summer Meeting, The Occasional Workshop in Environmental and Resource Economics, The Research Frontiers in the Economics of Climate Change Workshop, The University of Arizona, The University of Miami, and Iowa State University. Funding from The University of Arizona GPSC Research Grant is gratefully acknowledged.

Integrated assessment models (IAMs) are economists' primary tool for analyzing the optimal carbon tax. IAMs are macroeconomic models linked to a climate module by the "damage function," which translates rising temperature into economic losses.¹ The world's true underlying damage function is complex and characterized by deep uncertainties stemming from a lack of knowledge of how warming will affect natural and economic systems. In order to develop tractable models in the face of these unknowns, integrated assessment modelers have been forced to pin down the damage function with strong assumptions. Alongside the discount rate, the damage function is one of the most contentious feature of IAMs. In fact, some economists have suggested abandoning quantitative integrated assessment because of the alleged arbitrariness of the damage assumptions underpinning it (Pindyck, 2013).

Rather than abandoning the quantitative integrated assessment agenda that has proved useful for analysis and policymaking, I integrate uncertainty and even skepticism about damages into the benchmark DICE integrated assessment model. I advance methodology by incorporating uncertainty and Bayesian learning about the damage function. I also implement robust control, a macroeconomic technique, to account for concerns from economists and scientists that the damage functional form is misspecified. Using these methodological advances, I aim to ascertain the policy and welfare implications of: (1) accounting for parametric uncertainty over the damage function calibration, (2) including endogenous damage learning in a manner that closely matches how real world modelers update damage functions, (3) allowing the policymaker to distrust her model via the use of robust control, and (4) interacting learning with the policymaker's distrust of her own model.

Similar to approaches in the savings literature (e.g., Gollier, 2004), I decompose the optimal carbon tax into different channels representing the policy implications of uncertainty and learning, and I present a novel channel representing concerns about model misspecification. I find that subjective uncertainty over how the damage function is calibrated generally decreases the optimal carbon tax, however concern for model misspecification can flip the sign of this effect. The other uncertainty channels and the misspecification concern channel are analogous to precautionary savings and insurance motives in the savings literature. These channels tend to increase the optimal carbon tax at an increasing rate over time as the Earth warms and future damages appear more variable.

Even when the policymaker has misspecified the damage function, updating its calibration can lead to substantial ex-post welfare gains over a wide range of time horizons compared to

¹This is specific to the benchmark DICE model. Other IAMs, such as FUND, may have explicit impacts of CO₂.

not updating damage beliefs. Using robust control to guard against model misspecification can also improve welfare; however its benefits are realized far in the future, so net welfare gains relative to only acknowledging the uncertain damage calibration may not be achieved for decades or centuries.

Damage functions have garnered substantial criticism that may lead a policymaker to distrust damage functions in IAMs. Some economists have noted that IAMs do not account for a variety of potential damage channels (Howard, 2014), that the studies used to calibrate the damage function underestimate the actual impact of warming (Hanemann, 2008; Howard and Sterner, 2014), and that integrated assessment modelers have imposed strong, arbitrary functional form assumptions (Pindyck, 2013). Recent work analyzing the sensitivity of the optimal carbon tax to both the functional form and calibration of the damage function has demonstrated that errors in the damage function have non-trivial policy impacts (Stanton et al., 2009; Kopp et al., 2012; Weitzman, 2012).² Addressing these issues is critical since IAMs have become a key component of determining environmental regulation. For instance, a suite of IAMs is currently being used by the U.S. government to price greenhouse gas emissions in cost-benefit analyses of federal policies (Greenstone et al., 2013).

The weaknesses of current damage functions and uncertainty about the economic consequences of additional warming indicate that damage uncertainty should be explicitly captured in modeling. Damage uncertainty has been included only recently in IAMs (Croston and Traeger, 2013, 2014), however the recursive IAM literature, which uses dynamic programming approaches to better represent decision making under uncertainty, is relatively new. It has mostly focused on policy and learning when the sensitivity of the climate system to CO₂ is uncertain (Kelly and Kolstad, 1999; Leach, 2007; Kelly and Tan, 2015; Fitzpatrick and Kelly, 2016; Lemoine and Rudik, 2017), when there are tipping points or irreversibilities (Lemoine and Traeger, 2014; Cai et al., 2015a,b; Lontzek et al., 2015; Lemoine and Traeger, 2016a,b), or to examine the implications of solar geoengineering (Heutel et al., 2015, 2016).³

Robust control has also been applied recently to examine the effects of model misspecification concerns in environmental and integrated assessment contexts (Roseta-Palma and Xepapadeas, 2004; Athanassoglou and Xepapadeas, 2012; Anderson et al., 2014; Temzelides

²Economists often analyze warming risk, captured by climate sensitivity uncertainty. Damage uncertainty is less studied, but there is evidence that it has greater policy implications than warming uncertainty (Lemoine and McJeon, 2013).

³The recursive IAM literature expands on an expansive set of previous works that approximates uncertainty using Monte Carlo methods (Hope, 2006; Stern, 2006; Nordhaus, 2008; Ackerman et al., 2010; Kopp et al., 2012). Lemoine and Rudik (2017) provide a discussion and theoretical examples showing that Monte Carlo analyses are not equivalent to a dynamic programming approach.

et al., 2014). I build upon these strands of existing work by realistically capturing uncertainty, learning, and misspecification concerns in a recursive integrated assessment setting. To perform this analysis I integrate the benchmark DICE model into a dynamic programming setting that trades off near-term costs of emissions reductions with the future benefits of lower temperature via reduced damages in four different frameworks.

In order to develop these frameworks I must introduce additional state variables for the policymaker's beliefs and the stochastic realizations of damages. I exploit sparse grid methods which provide an accurate and computationally cheap way to substantially increase the number of states in dynamic stochastic IAMs. I use the original methodology of Smolyak (1963), and advances by Judd et al. (2014), to construct a sparse collocation grid for value function approximation. This methodology results in computational complexity increasing only *polynomially* in the number of states rather than exponentially (Winschel and Kratzig, 2010). This method has been applied recently in the macroeconomics literature to compute equilibria in overlapping generations models (Krueger and Kubler, 2004) and to solve stochastic growth models (Gonzalez and Rojas, 2009). This affords major increases in computational efficiency for high dimensional problems. This work appears to be the first application of this powerful computational technique in the environmental literature, which opens up new avenues for research in climate modeling.⁴

The paper is organized as follows. Section I provides more background on the construction of the DICE damage function and an overview of the DICE model and all four frameworks. Section II describes how to decompose the optimal carbon tax under uncertainty into several different channels. Section III reports results on how the different frameworks affect optimal policy and ex-post welfare when the damage function may be misspecified. Section IV concludes. The appendix fully describes the dynamic stochastic DICE model, contains additional information on robust control, details the computational methodology, and provides an error and sensitivity analysis.

I Damages and the Integrated Assessment Modeling Framework

I begin by giving a brief overview of the DICE damage function and its calibration. Next, I describe the DICE model as a whole and then describe the four frameworks I use to

⁴Lemoine and Rudik (2017) have subsequently used sparse grid methods and have called for applications of frontier numerical methods in the integrated assessment literature.

investigate damage uncertainty, learning, and misspecification concerns through the use of robust control.

I.I The DICE Damage Function Its Criticisms

In the benchmark DICE model, damages $D(T_t^s)$, multiplicatively reduce the level of time t output as a function of the time t surface temperature, T_t^s .⁵ The damage function specific to the DICE model is given by:

$$D(T_t^s) = d_1 [T_t^s]^{d_2},$$

where d_2 is set to 2 and d_1 is calibrated to estimates of damages from specific levels of warming over preindustrial levels.

Historically, the damage function has been calibrated in a bottom-up fashion by developing damage functions specific to different sectors such as health or agriculture, and then aggregating them up to a global damage function. The most recent version of DICE is directly calibrated to a set of monetized damage estimates contained in Tol (2009), with an upward adjustment of 25 percent to account for impacts that are more difficult to estimate such as human conflict (Nordhaus and Sztorc, 2013). Even with this change in the damage calibration procedure, it does not include damage estimates from a burgeoning strand of literature that has recently explored many different economic consequences of warming.⁶

Although the quadratic functional form is calibrated to data, the functional form itself is not. The quadratic functional form is a critical assumption and has been claimed to lack economic or scientific grounding (Stanton et al., 2009; Pindyck, 2013). Some scientists and economists have gone so far as to say the damage function is far removed from evidence in the natural sciences on catastrophes and impacts (Ackerman and Stanton, 2012; Lenton and Ciscar, 2012). Yet some assumption on its functional form is necessary because optimizing climate policy requires us to be able to quantify the benefits of abating emissions: the avoided damages. Since there is a lack of theory and knowledge to tell us what the functional form should be (Pindyck, 2013), modelers have resorted to imposing their own choice of functional forms.

⁵Although not investigated here, there are models that incorporate damages directly on capital (Kopp et al., 2012) or utility (Sterner and Persson, 2008; Barrage, 2016). Sterner and Persson (2008) explicitly include non-market environmental services in the utility function with limited substitutability, which can be degraded by warming. There is also empirical evidence that climate damages affect the growth rate of output, not just the level (Dell et al., 2012), and this may have substantial policy impacts (Moore and Diaz, 2015).

⁶See Hsiang (2016) for references to several recent examples.




Figure 1: A schematic of the DICE model with learning. The economy uses capital, labor and technology to produce output. Emissions from factor production flow into the carbon system. Additional atmospheric carbon dioxide increases temperature via radiative forcing. Higher surface temperature increases damage to output. A learning policymaker uses her observation of temperature and damages to refine her beliefs.

I.II Four Frameworks for Damage Uncertainty and The Dynamic Programming DICE Model

In light of the damage function’s critiques, I develop four different frameworks for investigating the policy and welfare implications of damage uncertainty. The “uncertainty framework” has the policymaker applying a distribution over d_1 and treating it as subjectively uncertain. The “learning framework” allows the policymaker to learn and refine her beliefs about d_1 over time. The “robust framework” has the policymaker being uncertain about the value of d_1 , and also applies robust control to capture concerns that the damage function in her model is misspecified. Finally, the “robust control and learning (RC+L) framework” has the policymaker combine updating her beliefs about d_1 with the use of robust control.

I investigate damage uncertainty within a dynamic programming version of the benchmark DICE model. The dynamic programming model (and the original DICE model) is finite-horizon Ramsey-Cass-Koopmans growth model coupled to a climate module with a ten year timestep.⁷ The state space is composed of capital; the atmospheric, upper ocean,

⁷Other economists have developed versions of the DICE model with a reduced state space (Traeger, 2014), or an annual timestep (Cai et al., 2015b). Lemoine and Rudik (2017) also use the same base model as this



Figure 2: The timeline for each period. The policymaker begins by observing the state of the world, then she updates her beliefs, and finally optimizes her policy for the current period.

and lower ocean CO_2 stocks; surface and ocean temperatures; a state to capture stochastic damages; and a state to capture a stochastically evolving parameter governing the policymaker's beliefs about d_1 .⁸

Figure 1 displays a schematic of the model. The model begins in 2005 and ends in 2605 with a given terminal value function. Each period, the economy begins with an existing level of capital, labor and technology. These are combined in a production function to generate output. The resulting output can be used in three ways: investment to increase the future stock of capital, abatement to reduce industrial CO_2 emissions caused by factor production, or consumption to increase flow utility. The policymaker's objective is to maximize her discounted expected stream of utility from consumption. When deciding on the optimal way to divide output between abatement, consumption, and investment, the policymaker uses her current beliefs about damages, which are given by her distribution over the damage calibration d_1 .

Emissions directly enter the atmospheric stock of CO_2 in the carbon system. CO_2 flows between the atmosphere and two other CO_2 reservoirs: the upper ocean and biosphere, and the lower ocean. Higher levels of atmospheric CO_2 increase radiative forcing, a measure of how much heat is trapped by greenhouse gases. Greater radiative forcing in turn raises surface temperature in the climate system. As the Earth's surface warms, it transfers thermal energy to the ocean, which increases ocean temperature. Higher surface temperature causes more damage, and reduces output. If the policymaker learns over time, she uses observations of surface temperature and damages in a given period to update her beliefs about d_1 .

The timing of the model is shown in Figure 2. The period begins and then the policymaker observes the state of the world. Damages are unknown prior to the beginning of the period since there is subjective uncertainty over the damage calibration and objective stochasticity

paper.

⁸As will be shown later, the shape parameter of her belief distribution evolves exogenously over time so it does not need to be kept as an explicit state.

in the realization of damages. After the policymaker observes the current state, she uses Bayes' Law to update her distribution over the damage calibration. Next, with her new beliefs, she selects levels of abatement, consumption, and investment that maximize her expected welfare. Finally, the world transitions to the next period. The policymaker's problem can be represented by a Bellman equation,

$$V_t(\mathbf{S}_t) = \max_{C_t, \alpha_t} \{U(C_t, L_t) + \beta E[V_{t+1}(\mathbf{S}_{t+1})]\}. \quad (1)$$

subject to:

$$\mathbf{S}_{t+1} = f(\mathbf{S}_t, C_t, \alpha_t)$$

where \mathbf{S}_t is time t state vector, α_t is abatement, C_t is consumption, L_t is labor, $U(C_t, L_t)$ is her flow utility function, $V_t(\mathbf{S}_t)$ is the policymaker's time t value function, and f is the set of state transition equations. Investment is determined by the residual output after consumption and abatement. In each period, the policymaker maximizes the sum of her current flow utility and her discounted expected continuation value where she takes expectations using her time t beliefs. The appendix contains the more detailed, full representation of the policymaker's problem.

I.III The Damage Calibration and Learning

In all four frameworks, the policymaker assigns the damage calibration a lognormal prior distribution, $d_1 \sim \log\mathcal{N}(\mu_t, \Sigma_t)$. In the frameworks with learning, the policymaker updates this distribution each period.⁹ I use a lognormal prior for a combination of tractability and to ensure the policymaker believes that higher temperatures lead to higher damages as in the existing literature on damage uncertainty (Crost and Traeger, 2013, 2014). Her immediate learning of the value of d_1 is hindered by a sequence of independent and identically distributed random shocks to damages, $w_t \sim \log\mathcal{N}(\mu_w, \sigma_w^2)$. These shocks capture random variation in how warming affects factor production through, for example, random variation in the frequency and magnitude of droughts or cyclones. The shock enters the damage

⁹The updating of solely d_1 and not the functional form approximates the learning process of real world integrated assessment modelers who have typically maintained the same damage function in updated versions of their models, but refined the calibration at points in time. For example, in the 1999 Regional Integrated Climate-Economy (RICE) model, low levels of warming increased output (Nordhaus, 2008), but in future model vintages, the damage function was updated so positive warming strictly reduces output. All models retained a quadratic functional form.

function multiplicatively,

$$D(T_t^s, w_t) = d_1 [T_t^s]^2 w_t,$$

so that the damage function also has a lognormal distribution: $D(T_t^s, w_t) \sim \log\mathcal{N}([T_t^s]^2 + \mu_t, \Sigma_t + \sigma_w^2)$. Damages reduce the level of gross output as in the benchmark DICE model (Nordhaus, 2008),

$$Y_t^n = \frac{Y_t^g}{1 + d_1 [T_t^s]^2 w_t}.$$

Rearranging the equation and taking the natural logarithm of both sides yields another equation such that observed variables are on one side, and unobserved variables are on the other,

$$\log\left(\frac{Y_t^g}{Y_t^n} - 1\right) - \log([T_t^s]^2) = \log(d_1) + \log(w_t).$$

The policymaker observes the stochastically evolving level of output net of damages, Y_t^n , and she infers gross output before damages, Y_t^g , using the production function, the level of technology, and the capital and labor stocks.¹⁰ This comprises all the variables on the left hand side. On the other side of the equation, $\log(d_1)$ and $\log(w_t)$ are normally distributed random variables that are not directly observed.

Relabel the expression on the left hand side Q_t , where now $Q_t \sim \mathcal{N}(\mu_t + \mu_w, \Sigma_t + \sigma_w^2)$. With each observation of the random variable Q_t equal to some realized value q_t at time t , the learning policymaker updates the parameters of her prior according to Bayes' Law,

$$\mu_t = \frac{\Sigma_{t-1}(q_t - \mu_w) + \sigma_w^2 \mu_{t-1}}{\Sigma_{t-1} + \sigma_w^2} \tag{2}$$

$$\Sigma_t = \frac{\Sigma_{t-1} \sigma_w^2}{\Sigma_{t-1} + \sigma_w^2}. \tag{3}$$

Σ_t , the shape parameter of the lognormal posterior for d_1 , declines monotonically and deterministically over time.¹¹ Since it is a function of only the previous period's shape parameter and the shape parameter of the shock distribution, there is no active learning channel in this model.¹² μ_t , the location parameter, is a weighted average of its previous value, and the

¹⁰This is an approximation to the real world learning process. In reality we may not know the global production function nor stock of capital; however, researchers have been estimating damages over time in different sectors using observations of temperature fluctuations and measures of production. See Hsiang (2016) for examples of recent research.

¹¹This is the variance parameter of the underlying normal distribution, and μ_t is the corresponding mean.

¹²There is only the possibility for active learning if the policymaker can amplify the signal relative to the noise through her control variables. In unknown climate feedback settings, this is done through allowing




Figure 3: The probability density function of the coefficient’s prior distribution in 2005 (left), and the probability density function for next period’s damages at 1°C, 2°C, and 3°C of warming given year 2005 beliefs (right).

realization of $q_{t+1} - \mu_\omega$, a noisy signal of the underlying value of d_1 .

The policymaker’s expectation of d_1 in the initial period is equal to the DICE calibration: $\mu_0 = 0.0028388$. The subjective variance of d_1 is set to 0.13^2 as in Nordhaus (2008). The probability density function the coefficient d_1 is shown in left panel of Figure 3.¹³ The mean of the lognormal shock is set to 1 so that the average shock has no effect on damages. I calibrate the variance of the shock so that the variance of $d_1 w_t$ is double that of d_1 alone, which is the range of damage variances used in the previous literature (Nordhaus, 2008; Crost and Traeger, 2013, 2014). The appendix contains a sensitivity analysis of the shock variance. The right panel of Figure 3 displays the subjective distribution of damages at 1°C, 2°C, and 3°C using the year 2005 prior over d_1 . The 1°C distribution is highly peaked at near-zero damages, and the right tail decays rapidly. The 2°C and 3°C distributions peak at slightly higher damages, but are characterized by right tails that decay more slowly and assign greater weight to higher damage outcomes.

additional emissions to accelerate warming (e.g., Kelly and Tan, 2015; Hwang et al., 2014) so the feedback signal swamps the temperature noise. Here, the shock enters multiplicatively with the unknown damage coefficient and thus, the noise and signal are equally amplified which provides no additional information. See Karp and Traeger (2013) for additional information on when active learning is possible.

¹³The densities in both panels integrate to 1.

I.IV Robust Control

Ideally, skepticism about damage functions should be built into an IAM by modeling a policymaker who does not completely trust that the model is a precise representation of reality. The policymaker instead would believe that her model is only an approximation to the real world. In this case, the policymaker would recognize that the optimal policy she obtains from her approximating model will almost surely not be optimal if taken and applied in real world policy. So instead of striving to develop an optimal policy using a model that is almost surely incorrect, the policymaker can instead try to find policy rules that are robust to unknown, and in the short-term unlearnable, errors in her model.¹⁴

To capture concerns about damage function misspecifications, I model a policymaker who uses robust control techniques in order to find policies that perform well even when the damage function inside her IAM may be incorrect (Hansen and Sargent, 2007, 2008). The policymaker begins with her approximating model, for example the standard quadratic damage function, and then selects her policy by finding one that fares well over a set potential models that are close in terms of the Kullback-Leibler divergence. The true model likely resides in this set; however, because the models in this set are all relatively similar to her own, they are difficult to statistically distinguish from one another over the time frame of her policymaking. During policy optimization, alternative models more similar to her approximating model are given greater weight, while those that are more different get less weight. By taking this approach the policymaker is recognizing that her approximating damage function is probably close to the true damage function, but not exactly right, and she wishes to guard against unknown errors.¹⁵

Robust control induces robust decision rules by distorting state transition distributions (Hansen and Sargent, 2008). The distortions reduce her perceived expected continuation value and induce her to select policies that are designed for lower welfare worlds. Importantly, in simulations the distortions only affect the policymaker's decision rule, and not the actual transition distributions.¹⁶ The policymaker maximizes her objective while accounting for

¹⁴Implicitly the four frameworks are taking different approaches to estimating an unknown data generating process for damages. In the uncertainty framework a quadratic damage function is estimated once before policymaking begins. The learning framework allows for periodic updates to the quadratic model, the robust framework considers the possibility of other approximations to the damage function, and the framework combining robust control and learning allows for the possibility of other damage functions with continual re-estimation of the calibrated coefficient on the approximating damage function.

¹⁵This is in contrast to maxmin expected utility which maximizes welfare subject to the worst-case model in some set (Gilboa and Schmeidler, 1989).

¹⁶Hansen and Sargent (2008) also demonstrate that the distortions can be represented as the optimal

potential distortions, weighted by their Kullback-Leibler divergence. The distortions the policymaker faces are implemented via state-dependent shocks to the transition distributions, so that the distortions can be temporally linked and persist over time. This allows the distortions to represent real misspecifications to the policymaker’s model. The distortions are implemented in a non-linear dynamic programming framework by replacing the continuation value in the policymaker’s Bellman equation with a risk sensitivity operator T^1 (Hansen and Sargent, 2007),

$$T^1(\beta V_{t+1}(\mathbf{S}_{t+1})|\theta_1) = -\theta_1 \log \left(E_{\mathbf{S}_{t+1}} \left[\exp \left(-\frac{\beta V_{t+1}(\mathbf{S}_{t+1})}{\theta_1} \right) \right] \right),$$

where θ_1 is a penalty parameter, described below, that calibrates how strongly the policymaker guards against these misspecifications, and the expectation operator is over next period’s state, but not over d_1 . This results in a new Bellman equation,

$$V_t(\mathbf{S}_t) = \max_{C_t, \alpha_t} \{U(C_t, L_t) + E_{d_1} [T^1(\beta V_{t+1}(\mathbf{S}_{t+1})|\theta_1)]\}. \quad (4)$$

subject to:

$$\mathbf{S}_{t+1} = f(\mathbf{S}_t, C_t, \alpha_t)$$

where the expectation acting outside the risk sensitivity operator is over d_1 .

For any $\theta_1 \in (0, \infty)$, the risk sensitivity operator distorts the density of continuation values towards worse outcomes than if it was the usual expected continuation value. As $\theta_1 \rightarrow \infty$, the risk sensitivity operator reduces to the usual subjective expected utility framework.¹⁷ There typically exists some point $\bar{\theta}_1$, where the problem “breaks down” for $\theta_1 < \bar{\theta}_1$. I calibrate θ_1 by selecting close to the breakdown point that would induce high levels of concern for model misspecification stemming from the points outlined in Section I.I.¹⁸ For

selection of a transition probability measure by an evil agent whose objective is to minimize the payoff of our policymaker. The evil agent’s selection is subject to a penalty proportional to the Kullback-Leibler divergence of its distorted transition probability measure relative to the approximating model’s transition probability measure.

¹⁷Consider a simple case where the continuation value, V_{t+1}^r , is distributed $N(1, 1)$. The expected continuation value is just 1. However, when we apply the risk sensitivity operator, the term inside the expectation is now lognormally distributed and the expected continuation value becomes $1 - \frac{1}{2\theta_1}$. Clearly, as $\theta_1 \rightarrow \infty$, we are back in the subjective expected utility world. As θ_1 decreases, the expected continuation value declines and the policymaker acts as if she is facing worse futures. This induces her to select policies that better guard against misspecifications which reduce her future welfare.

¹⁸Hansen and Sargent (2008) demonstrate how to calibrate θ_1 using detection error probabilities in a linear control setting when the distortion to the transitions can be solved for. Athanassoglou and Xepapadeas

both frameworks the penalty parameter is set to $\theta_1 = 200000$. In Section III.II I show how varying the concern for model misspecification alters policy.

II The Policy Effects of Uncertainty and Concern for Model Misspecification

The time t optimal carbon tax is conventionally defined as the shadow value of time t emissions relative to capital. Time t emissions affects time $t + 1$ atmospheric CO₂, surface temperature, and the fraction of output remaining after damages ($\mathcal{L}_{t+1} = \frac{1}{1+d_1 [T_t^s]^2 w_t}$). From here on I'll call this fractional net output.¹⁹ To keep the equations as concise as possible without losing intuition I focus only on the the temperature and fractional net output terms. The omitted CO₂ terms will be identical to the temperature terms but with temperature replaced by CO₂. The optimal carbon tax is then,²⁰

$$\text{Tax}_t = E \left[\frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \frac{\partial T_{t+1}^s}{\partial e_t} + \frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] \bigg/ E \left[\frac{\partial V_{t+1}}{\partial K_{t+1}} \right].$$

Without loss of generality, consider the optimal carbon tax for a policymaker utilizing robust control. The optimal carbon tax expression is then,

$$\text{Tax}_t = \frac{E \left[\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right) \left(\frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \frac{\partial T_{t+1}^s}{\partial e_t} + \frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right) \right]}{E_{\mathbf{S}_{t+1}} \left[\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right) \right]} \bigg/ E \left[\frac{\partial V_{t+1}}{\partial K_{t+1}} \right],$$

After taking the derivative of the risk sensitivity operator, the expectation is over both d_1 and \mathbf{S}_{t+1} in the numerator. The application of robust control introduces a new term, $\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right)$, which when normalized by its expectation, corresponds to the worst-case distortion to the perceived transition density induced by the risk sensitivity operator (Hansen and Sargent, 2007).

(2012) formally derive a closed form solution for the worst case misspecification exploiting a linear-quadratic model. The DICE model is highly non-linear and does not admit closed form solutions for these approaches to penalty parameter calibration.

¹⁹I use fractional net output as a state instead of damages to obtain a bounded domain, $\mathcal{L}_{t+1} \in [0, 1]$.

²⁰Emissions at time t only enter the atmosphere after one period in DICE. The capital shadow value can be thought of as coming from the policymaker's first-order condition for consumption where she trades off current utility against future capital. Hence, this is equivalent to normalizing by the marginal utility of consumption in time t , modulo a discount factor.

After dropping value function partial derivative with respect to capital for clarity, the equation can be rearranged to recover the conventional carbon tax expression (first line) and an additively separable adjustment for robust control (second line),

$$\begin{aligned} \text{Tax}_t \propto & E \left[\frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \right] \frac{\partial T_{t+1}^s}{\partial e_t} + E \left[\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \right] E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] + \text{cov} \left(\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}, \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right) \\ & + \text{cov} \left(\frac{-\partial V_{t+1}}{\partial \mathbf{S}_{t+1}} \frac{\partial \mathbf{S}_{t+1}}{\partial e_t}, \frac{\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right)}{E_{\mathbf{S}_{t+1}} \left[\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right) \right]} \right). \end{aligned}$$

And finally, a second-order Taylor expansion of the value function partial derivative terms on the first line around the expected time $t + 1$ state vector $\zeta := E[\mathbf{S}_{t+1}]$ results in the following expression, which I will group into six different channels,²¹

$$\begin{aligned} \text{Tax}_t \propto & \left. \frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \right|_{ce} \frac{\partial T_{t+1}^s}{\partial e_t} + \left. \frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \right|_{ce} \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \Big|_{ce} \tag{5} \\ & + \left(\left. \frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \right|_{\zeta} - \left. \frac{-\partial V_{t+1}}{\partial T_{t+1}^s} \right|_{ce} \right) \frac{\partial T_{t+1}^s}{\partial e_t} + \left(\left. \frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \right|_{\zeta} E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] - \left. \frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}} \right|_{ce} \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \Big|_{ce} \right) \\ & + \frac{1}{2} \frac{-\partial^3 V_{t+1}}{\partial T_{t+1}^s \mathcal{L}_{t+1}^2} \frac{\partial T_{t+1}^s}{\partial e_t} \text{var}(\mathcal{L}_{t+1}) + \frac{1}{2} \frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}^3} E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] \text{var}(\mathcal{L}_{t+1}) \\ & + \frac{-\partial^3 V_{t+1}}{\partial T_{t+1}^s \partial \mathcal{L}_{t+1} \partial \mu_{t+1}} \frac{\partial T_{t+1}^s}{\partial e_t} \text{cov}(\mathcal{L}_{t+1}, \mu_{t+1}) + \frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}^2 \partial \mu_{t+1}} E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] \text{cov}(\mathcal{L}_{t+1}, \mu_{t+1}) \\ & + \frac{1}{2} \frac{-\partial^3 V_{t+1}}{\partial T_{t+1}^s \partial \mu_{t+1}^2} \frac{\partial T_{t+1}^s}{\partial e_t} \text{var}(\mu_{t+1}) + \frac{1}{2} \frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1} \partial \mu_{t+1}^2} E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] \text{var}(\mu_{t+1}) \\ & + \text{cov} \left(\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}, \frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right) + \text{cov} \left(\frac{-\partial V_{t+1}}{\partial \mathbf{S}_{t+1}} \frac{\partial \mathbf{S}_{t+1}}{\partial e_t}, \frac{\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right)}{E_{\mathbf{S}_{t+1}} \left[\exp \left(\frac{-\beta V_{t+1}}{\theta_1} \right) \right]} \right), \end{aligned}$$

where $|_{ce}$ indicates certainty: the term is evaluated as if the policymaker is in a deterministic world with d_1 equal to the mean of her time t beliefs. For the subsequent analysis of channels, note that fractional net output is decreasing in emissions ($\partial \mathcal{L}_{t+1} / \partial e_t < 0$) and temperature is increasing in emissions ($\partial T_{t+1} / \partial e_t > 0$).

The first channel is the *certainty* tax which is composed of the terms on the first line on the right hand side of equation (5). This is the tax the policymaker would set in a deterministic, perfect-information world if she happened to be at her current state at time t

²¹Lemoine and Rudik (2017) similarly demonstrate how to use Taylor expansions of the value function to gain insight into how uncertainty over the climate's sensitivity to emissions drives optimal policy.

and the damage calibration matched her expectation.

The second channel is the *uncertainty adjustment*, which is composed of terms on the second line. This channel alters the certainty tax so that it correctly accounts for uncertainty in the marginal effect of emissions on time $t+1$ expected welfare due to subjective uncertainty over the damage calibration's value and objective uncertainty over the damage shock in time $t + 1$.

The third channel is the *precautionary abatement* motive, which is composed of terms on the third and fourth lines. The third derivative of utility corresponds to how uncertainty about future consumption affects the agent's contemporaneous saving decision (Leland, 1968; Dreze and Modigliani, 1972; Kimball, 1990). If uncertainty about future consumption increases contemporaneous savings, the agent is said to exhibit prudence. In a climate-economy setting, abatement is a form of environmental savings: increasing current abatement means that the policymaker forgoes a sure consumption payoff now for increased consumption later due to lower future damages. The last term on the third line captures this most clearly. If the third derivative of her continuation value with respect to fractional net output is positive, then,

$$\frac{1}{2} \frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}^3} E \left[\frac{\partial \mathcal{L}_{t+1}}{\partial e_t} \right] \text{var}(\mathcal{L}_{t+1}) > 0,$$

and the policymaker increases her abatement. The precautionary motive also scales in size with the amount of subjective variability in future output. The previous two terms on the third line capture similar precautionary abatement motives in the face of uncertain output. If these terms increase abatement then the agent would be called cross-prudent (Gollier, 2010). The policymaker is cross-prudent if she would prefer to have a mean-zero risk attached to fractional net output when CO_2 or temperature are lower rather than when they are higher.

The fourth line also capture precautionary abatement motives because of cross-prudence, but here, instead of increasing abatement because of future variability in fractional net output, the policymaker abates because of co-variability between fractional net output and her beliefs about the damage calibration. This covariance is likely negative since lower than expected future output is the signal that the policymaker would need to receive to revise her future beliefs upward. This covariance matters because when future beliefs about damages are uncertain, then future output and future consumption appear to be even more variable. Whatever information the policymaker receives about output in the future feeds back onto her future payoffs through her expectations. If this covariance is negative, it tends to increase precautionary abatement.

The fourth channel is the *signal smoothing* motive, which is composed of the terms on the fifth line. This channel captures effects of learning on the optimal tax. The second derivative of the value function with respect to the location parameter, $-\partial^2 V_{t+1}/\partial\mu_{t+1}^2$, captures how well the policymaker can use new information to smooth welfare over possible values of d_1 . Since a higher location parameter for the policymaker's beliefs strictly increases her expectation about d_1 , and if $-\partial^2 V_{t+1}/\partial\mu_{t+1}^2 > 0$ as we may expect, then the marginal welfare cost of a higher damage calibration belief is increasing and convex. If there is not much curvature in her beliefs and $\partial^2 V_{t+1}/\partial\mu_{t+1}^2$ is small, then she is able to smooth welfare effectively, but the larger $\partial^2 V_{t+1}/\partial\mu_{t+1}^2$ is in magnitude, the less she can smooth welfare, and the more costly a bad signal of d_1 becomes.

The triple derivatives then indicate how a marginal increase in CO₂, temperature or output net of damages affects the policymaker's ability to smooth welfare in response to new information. Indeed, we may expect additional CO₂ or temperature decrease her ability to smooth welfare ($-\partial^3 V_{t+1}/\partial M_{t+1}^{atm} \partial\mu_{t+1}^2 > 0$, $-\partial^3 V_{t+1}/\partial T_{t+1}^s \partial\mu_{t+1}^2 > 0$), and that having a greater fraction of her gross output would increase her ability to smooth welfare ($-\partial^3 V_{t+1}/\partial \mathcal{L}_{t+1} \partial\mu_{t+1}^2 < 0$). Additional temperature or CO₂ results in greater damage and less output to be able to use towards capital investment or abatement, while additional output has the opposite effect. In these cases, additional variability in future beliefs magnifies this effect and increases the optimal tax.

The fifth channel is the *output insurance* channel, which is captured by the first term on the last line. Output insurance reduces the optimal level of emissions if and only if the covariance is positive. The second term in the covariance captures the marginal reduction in fractional net output from emissions. The first term in the covariance captures the marginal welfare cost of less fractional net output. Similar to consumption-based asset pricing models, the policymaker cares about the covariance of returns with marginal welfare. If emissions reductions are most effective in preserving output when output is most valuable to welfare, then this channel reduces the optimal level of emissions. Intuitively, this seems like the most plausible case since emissions reduce output most greatly when d_1 is large, and large d_1 implies lower output and higher marginal welfare.

The final channel is the *misspecification insurance* channel, which is captured by the final term. The misspecification insurance channel reduces emissions if and only if the covariance is positive. The magnitude and direction of the misspecification insurance channel depends on how the marginal welfare cost of emissions covaries with the distortion to the transition density. If a larger distortion makes the marginal welfare cost of emissions appear larger as




Figure 4: The mean optimal carbon tax over 5000 simulations for each framework when the true damage function is quadratic.

intuition suggests, then misspecification insurance tends to reduce emissions.

III Results

First, I display the mean optimal carbon tax trajectories and the corresponding climate outcomes for each of the four frameworks. Next, I vary the robust control penalty parameter to examine how the policymaker’s concern for model misspecification affects policy. Then I decompose the carbon taxes into each channel outlined in Section II to determine how uncertainty matters for policy. Finally, I analyze ex-post welfare outcomes under the different frameworks when the policymaker may have misspecified the damage function in her model. In each simulation run I sample d_1 from its year 2005 prior distribution and I also sample a vector of annual shocks from the distribution of the stochastic damage shock ω_t .

Table 1: The average carbon tax in 2015 and 2105, peak CO₂ and temperature, and the 5th and 95th percentiles (in parentheses) for each framework.

	Framework			
	Uncertainty	Learning	Robust	RC+L
2105 Tax (\$/tCO ₂)	50 (48,52)	48 (27,77)	53 (51,55)	48 (27,77)
2205 Tax (\$/tCO ₂)	184 (166,196)	175 (97,276)	189 (170,201)	175 (96,275)
Peak CO ₂ (ppm)	704 (700,709)	735 (609,887)	691 (687,696)	735 (609,886)
Peak Temperature (°C)	3.67 (3.65,3.71)	3.79 (3.10,4.44)	3.60 (3.58,3.64)	3.79 (3.10,4.44)

III.I Optimal Carbon Tax Trajectories

Figure 4 displays the first century’s mean optimal carbon tax trajectories in \$/tCO₂ when the policymaker has correctly specified the damage functional form as quadratic.²² In the uncertainty framework, the policymaker begins her optimal tax in 2005 at \$7.77/tCO₂ and ramps it up to \$50.40/tCO₂ in 2105. When the policymaker is able to learn the damage calibration over time, she sets a lower initial carbon tax of \$7.50/tCO₂, which rises to \$48.21/tCO₂ at the end of the century. Applying robust control to guard against potential model misspecification tends to increase the tax when not updating beliefs about d_1 . In 2005 the robust control tax begins at \$8.57/tCO₂, 10 percent higher than the uncertainty framework. Over the first century, the robust control tax increases faster than the uncertainty tax in absolute terms and reaches \$53.48/tCO₂ at the end of the century. Applying robust control on top of learning has negligible effects on the optimal carbon tax.

Table 1 displays the mean, 5th and 95th percentile outcomes for the 2105 carbon tax, the 2205 carbon tax, and for peak CO₂ and warming since year 1900. Learning results in significant variability in realized carbon taxes since the policymaker adapts her policy to the noisy information she receives about each simulation’s specific d_1 . After 100 years the learning carbon tax is on average only \$2/tCO₂ less than the uncertainty tax, but in some cases it may be nearly 50 percent lower or 50 percent higher. This also holds true after 200 years. After 100 years the mean robust control carbon tax is 6 percent higher than the mean

²²Some papers report carbon taxes in terms of dollars per ton of carbon. The unit conversion is 11 tons of CO₂ to three tons of carbon.

uncertainty carbon tax, but after 200 years the relative difference declines to 3%.

When learning, the policymaker peaks CO₂ 4 percent higher on average, which results in a 3 percent higher peak warming compared to the uncertainty framework. Conversely, the more aggressive carbon tax in the robust control framework tends to peak CO₂ and warming lower than the uncertainty framework, both by 2 percent. There is also significant variability in climatic outcomes when learning. The 90 percent confidence interval for peak CO₂ is 278 ppm wide, more than half of the approximately 400 ppm CO₂ concentration today in the real world. The 90 percent confidence interval in peak warming is 0.34°C, which is approximately how much the Earth warmed between 1980 and 2000.

III.II The Effect of Concern for Model Misspecification

Figure 5 displays how changing the level of the policymaker's concern about model misspecification affects the initial optimal carbon tax. Smaller values of the penalty parameter indicate greater concern. The left panel of Figure 5 displays the optimal carbon tax as a function of the penalty parameter for the uncertainty framework. When the penalty parameter is sufficiently large, the year 2005 optimal carbon tax is effectively equal to the optimal carbon tax when not applying robust control. As the penalty parameter declines and we move to the left on the plot, the optimal carbon tax increases, and on this plot, peaks at \$8.58/tCO₂ when the penalty parameter is 200000. Decreasing it further illustrates how the problem begins to breakdown. There is a sharp decrease in the optimal carbon tax at a penalty parameter around 110000, and then for values that are only slightly smaller, the model fails to solve.²³

The right panel of Figure 5 displays the initial optimal carbon tax for the robust control and learning framework. Again, a sufficiently large penalty parameter leads the optimal carbon tax to be effectively equal to the learning framework without robust control. Decreasing the penalty parameter leads to a slight increase in the carbon tax before beginning a quicker decline once the penalty parameter goes below 200000. The effect of varying the concern for model misspecification is substantially smaller than without learning. The optimal carbon tax reaches a low point at a penalty parameter around 130000; however, the initial carbon tax never varies by more than \$0.01/tCO₂. Decreasing the penalty parameter to slightly below 130000 leads to a sharp increase in the optimal tax to \$8.49/tCO₂ (off the plot) and

²³Since the policymaker has an iso-elastic utility function with an elasticity of intertemporal substitution that is less than 1, her utility is negative and unbounded below. The problem breaks down when her continuation value is numerically equivalent to negative infinity.





Figure 5: The year 2005 optimal carbon taxes when using robust control in the uncertainty framework (left) and the learning framework (right). Note that each panel has a different y-axis scale and the x-axis is on a log10 scale.

then a complete breakdown of the problem.²⁴

III.III Decomposing the Optimal Carbon Tax

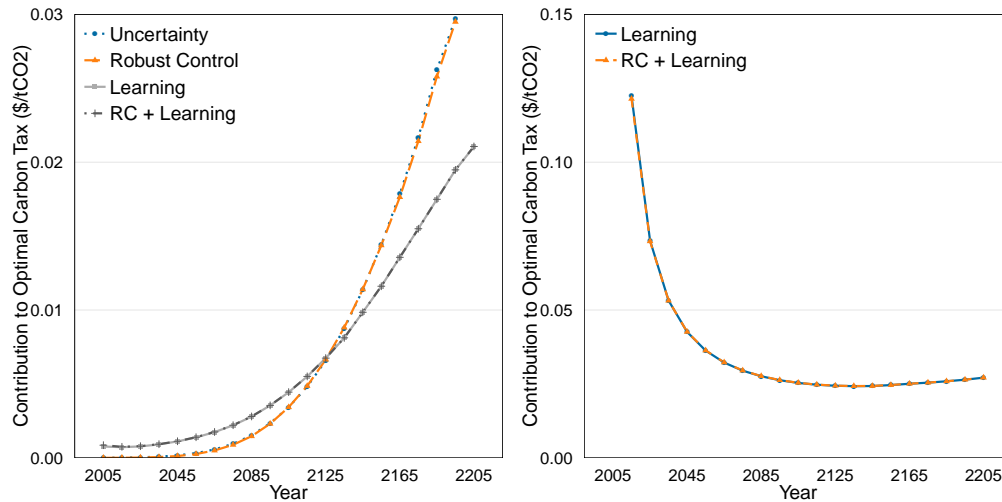
Figure 6 plots the five channels that capture uncertainty’s effect on the optimal carbon taxes plotted in Figure 4. Note that all panels have different scales to better display differences between the different frameworks. Panel (a) plots the uncertainty adjustment, the strongest of the channels. For all but the robust control framework, the uncertainty adjustment is negative and grows in magnitude over time. This negative adjustment is driven by the convexity of output in damages. As the Earth’s surface warms, it increases damage variability since surface temperature enters damages multiplicatively with the uncertain damage calibration and the stochastic damage shock. Because output is convex in damages, increased variability yields higher expected output and greater welfare. The policymaker’s optimal adjustment is then to reduce the carbon tax. For the uncertainty framework, the adjustment begins at nearly zero, but monotonically increases in size to $-\$2/\text{tCO}_2$ in 2205 as temperature increases

²⁴Selecting an alternative penalty parameter than the one used here for the learning framework has virtually no impact on the results.




(a) Uncertainty Adjustment

(b) Output Insurance



(c) Precautionary Abatement

(d) Signal Smoothing



(e) Misspecification Insurance

Figure 6: The five carbon tax channels. Note that each panel has a different y-axis scale.

and damages appear subjectively more variable to the policymaker.

In the two frameworks with learning the uncertainty adjustment is more strongly negative since increases in damage variability are magnified by the policymaker's anticipated variability in future beliefs. In the frameworks with learning, the uncertainty adjustment increases in size from approximately $-\$0.60/\text{tCO}_2$ to nearly $-\$10.00/\text{tCO}_2$ between 2005–2205.

The robust control framework is the only one with a positive uncertainty adjustment. The risk sensitivity operator, T^1 distorts the continuation value in a way that makes uncertainty over the damage calibration appear to result in lower expected future welfare, which induces an increase in the carbon tax. This effect increases over time from $\$0.80/\text{tCO}_2$ in 2005 to $\$2.90/\text{tCO}_2$ in 2145 before beginning to decline.

Panel (b) displays the effect of output insurance. As hypothesized, this channel increases the optimal carbon tax. This channel is near zero at first since surface temperature is low and there is not much variability in the marginal effect of emissions on fractional net output. After two centuries, the size of the channel grows to approximately $\$0.60/\text{tCO}_2$ when not learning. The two learning frameworks see a smaller amount of output insurance which peaks at only $\$0.25/\text{tCO}_2$. Learning resolves subjective variability in the effect of emissions on fractional net output, which offsets the increased variability caused by higher temperature.

Panel (c) plots the precautionary abatement motive, which is small over the entire time horizon. Initially precautionary abatement is effectively zero since variability in damages is negligible early on. Early precautionary abatement is higher when learning since the cross-prudent policymaker is additionally concerned about how variability in future damages covaries with uncertain future beliefs. Additional warming over time increases damage variability, leading to greater precautionary abatement for all frameworks. In the learning frameworks, this is again partially offset by the policymaker resolving her subjective uncertainty over the damage calibration. In all cases, the size of the precautionary abatement motive is small and contributes at most a few cents to the optimal carbon tax.

Panel (d) plots the signal smoothing channel, which is specific to the frameworks with learning. This channel increases the optimal carbon tax because emissions hinder the policymaker's ability to smooth welfare. The signal smoothing channel peaks at $\$0.13/\text{tCO}_2$ in 2005 and declines over time as she resolves uncertainty over d_1 and her future beliefs become less variable. With each update of her beliefs, more weight is placed on her prior and new information becomes less valuable. In turn, emissions become less costly.

Panel (e) plots the misspecification insurance channel. The channel always increases

the carbon tax, indicating that the distortion to her transition density makes emissions reductions appear especially valuable. This channel begins small but grows over time for both robust control frameworks. The marginal benefit of emissions reductions and the distortion to the transition density become more variable, which increases their covariance and the desire to insure against misspecification. The misspecification insurance channel grows slower when learning since resolving uncertainty over d_1 reduces variability in both terms in the covariance. This results in misspecification insurance only half as large in 2205 as when not learning.

III.IV Ex-Post Welfare when the Damage Function is Misspecified

I next investigate the relative welfare performance of the frameworks when facing true damage functions that may be different than the one in the policymaker’s model. I assume the policymaker has correctly specified the damage function as within the class of polynomials, and that her distribution on the coefficient of d_1 is also correct. However, I allow the degree of the polynomial to be misspecified so the function is either less convex with an exponent of 1.75, or more convex with an exponent of 2.25 or 2.5.

Since I am exploring outcomes when models are misspecified, I must simulate the model and calculate the welfare outcomes over a finite horizon. The frameworks’ value functions only yield the true ex-ante expected welfare when the model is specified correctly. This welfare analysis is also ex-post. Although an ex-ante analysis is intrinsically attractive, it requires assuming that economists and scientists have a well-grounded distribution for the degree of the damage function. I do not take a stand on what the distribution over the damage exponent should be and instead I calculate welfare for specific values for the degree of the damage function. An ex-ante analysis can be easily approximated by weighting the welfare outcomes for each of the damage functions examined here.

Figure 7 displays the balanced growth equivalent (BGE) gain in consumption for each of the frameworks relative to the uncertainty framework (Mirrlees and Stern, 1972; Jensen and Traeger, 2013; Lemoine and Traeger, 2014). The BGE is the per-capita consumption level which, when growing at some constant rate, would yield the same level of welfare as the the optimal policy of a given framework. The horizontal axis on each plot denotes four different horizons over which welfare is evaluated. Each plot is for a different exponent. The uncertainty framework is omitted since it is the baseline framework. An equivalent level of welfare as the uncertainty framework is denoted by the dashed line. Plots above the dashed




Figure 7: The balanced growth equivalent consumption gain of the learning (Learn, circle), robust control (RC, triangle), and robust control and learning (RC+L, square) frameworks relative to the uncertainty framework over four different time horizons and four different damage functional forms.

line indicate gains relative to the uncertainty framework and plots below the dashed line indicate losses.

If the true damage function exponent is 1.75 then the true damage function is less convex than the policymaker believes and damages will not rise as rapidly in the future. The robust control framework does worse than the uncertainty framework, incurring BGE losses of between 0.017 percent and 0.005 percent depending on the time horizon. The robust control framework guards against bad misspecifications by distorting future states to appear worse than the model indicates, but here, future states are actually better. Conversely, the frameworks with learning generally result in BGE gains. When learning, the BGE gain grows larger as the horizon is extended and peaks at 0.007 percent when considering a time horizon of 200 years. Although the policymaker has misspecified the damage function, learning does allow her to push expectations of future damages in the correct direction by attributing lower damage realizations to a lower damage calibration instead of a low damage exponent.

When the true damage function exponent is 2.0 and the model is correctly specified, the learning frameworks again do best with BGE gains of 0.001–0.008 percent. The robust control framework is incorrectly concerned for model misspecification, leading to a carbon tax that is too high and consumption that is too low. Robust control consequently incurs BGE losses over all the time horizons, but potentially up to 0.015 percent in the short term as the policymaker substitutes away from early consumption towards emissions abatement.

If the true damage function exponent is greater than 2.0, then the policymaker’s model underestimates the convexity of the damage function and how rapidly damages will increase as the Earth warms. When the exponent is 2.25, the learning frameworks again do best, with the robust control and learning framework performing slightly better than just learning alone. BGE gains under the learning frameworks reach up to 0.015 percent when the damage function exponent is 2.5. Without learning, robust control performs the worst in the short term when facing a more convex damage function. The benefits of increased abatement under robust control realize in the future and are spread out over many decades, so the policymaker needs a long time horizon to reap the benefits of forgone early consumption.

IV Conclusions

Most economists believe that the damage functions in IAMs are misspecified. Yet analyses of optimal climate policy have not investigated the effects of updating the damage function over time, or for adapting policy to take account of the widely noted concerns of damage function

misspecification. I fill this gap in the literature by comparing the performance of four policy frameworks with different degrees of damage assumptions. The uncertainty framework accounts for uncertainty in the calibration of the damage function while assuming a known functional form. The learning framework acknowledges that economists do learn over time and allows the policymaker to update her distribution over the uncertain damage function calibration within the model, assuming the same functional form. The robust framework borrows techniques from the macroeconomic literature to incorporate concerns that the damage function is misspecified in unknown ways. The robust policymaker constructs policies that guard against model misspecification. Last, the robust control and learning framework combines concern for model misspecification with updating of beliefs over the damage coefficient. For each of the frameworks, I analytically demonstrate how misspecification concerns, and uncertainty in damages and future beliefs feed into the optimal carbon tax.

I find that uncertainty over the damage calibration generally lowers the optimal carbon tax through the uncertainty adjustment. An analytical decomposition reveals that there are conflicting effects of uncertainty on the optimal carbon tax due to precautionary abatement, insurance against co-variability of the marginal cost of emissions and welfare, and variability in future beliefs. Concern for robustness tends increase the carbon tax through a channel that parallel insurance motives in the savings literature. If the damage function is misspecified as widely believed, robust control and (incorrect) learning can achieve higher welfare than accounting for parametric damage uncertainty alone. However, the welfare gains from utilizing robust control alone may not realize until decades or centuries into the future, raising potential concerns about inter-generational equity in bearing the costs and benefits of designing a more robust climate policy.

References

- Ackerman, F and Elizabeth A. Stanton (2012) “Climate Risks and Carbon Prices: Revising the Social Cost of Carbon,” *Economics : The Open-Access, Open-Assessment E-Journal*, Vol. 6, No. 10.
- Ackerman, Frank, Elizabeth A. Stanton, and Ramón Bueno (2010) “Fat Tails, Exponents, Extreme Uncertainty: Simulating Catastrophe in DICE,” *Ecological Economics*, Vol. 69, No. 8, pp. 1657–1665.
- Anderson, Evan W., William A. Brock, Lars Peter Hansen, and Alan H. Sanstad (2014) “Ro-

- bust Analytical and Computational Explorations of Coupled Economic-Climate Models with Carbon-Climate Response.”
- Athanassoglou, Stergios and Anastasios Xepapadeas (2012) “Pollution Control with Uncertain Stock Dynamics: When, and How, to be Precautious,” *Journal of Environmental Economics and Management*, Vol. 63, No. 3, pp. 304–320.
- Barrage, Lint (2016) “Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy.”
- Cai, Yongyang, Kenneth L Judd, Timothy M Lenton, Thomas S Lontzek, and Daiju Narita (2015a) “Environmental Tipping Points Significantly Affect the Cost-Benefit Assessment of Climate Policies,” *Proceedings of the National Academy of Sciences*, Vol. 112, No. 15, pp. 4606–4611.
- Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek (2015b) “The Social Cost of Carbon with Economic and Climate Risks,” pp. 1–58.
- Crost, Benjamin and Christian P. Traeger (2013) “Optimal Climate Policy: Uncertainty Versus Monte Carlo,” *Economics Letters*, Vol. 120, No. 3, pp. 552–558.
- (2014) “Optimal CO2 Mitigation Under Damage Risk Valuation,” *Nature Climate Change*, Vol. 4, No. 7, pp. 631–636.
- Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2012) “Temperature Shocks and Economic Growth: Evidence from the Last Half Century,” *American Economic Journal: Macroeconomics*, Vol. 4, No. 3, pp. 66–95.
- Dreze, Jacques H. and Franco Modigliani (1972) “Consumption Decisions Under Uncertainty,” *Journal of Economic Theory*, Vol. 5, No. 3, pp. 308–335.
- Fitzpatrick, Luke G and David L Kelly (2016) “Probabilistic Stabilization Targets,” *Journal of the Association of Environmental and Resource Economists*.
- Gilboa, I and D Schmeidler (1989) “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, Vol. 18, pp. 141–153.
- Gollier, Christian (2004) *The economics of Risk and Time*: MIT press.

- (2010) “Ecological Discounting,” *Journal of economic theory*, Vol. 145, No. 2, pp. 812–829.
- Gonzalez, A Gavilan and JA Rojas (2009) “Solving Portfolio Problems with the Smolyak-Parameterized Expectations Algorithm.”
- Greenstone, M., E. Kopits, and a. Wolverton (2013) “Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation,” *Review of Environmental Economics and Policy*, Vol. 7, No. 1, pp. 23–46.
- Hanemann, WM (2008) “What is the Economic Cost of Climate Change?,” *CUDARE Working Papers*.
- Hansen, Lars Peter and Thomas J. Sargent (2007) “Recursive Robust Estimation and Control Without Commitment,” *Journal of Economic Theory*, Vol. 136, No. 1, pp. 1–27.
- (2008) *Robustness*, Princeton, N.J.: Princeton University Press.
- Heutel, Garth, Juan Moreno Cruz, and Soheil Shayegh (2015) “Solar Geoengineering, Uncertainty, and the Price of Carbon.”
- (2016) “Climate Tipping Points and Solar Geoengineering,” *Journal of Economic Behavior & Organization*.
- Hope, Chris (2006) “The Marginal Impact of CO2 from PAGE2002: An Integrated Assessment Model Incorporating the IPCC’s Five Reasons for Concern,” *Integrated Assessment*, Vol. 6, pp. 19–56.
- Howard, Peter (2014) “Omitted Damages: What’s Missing from the Social Cost of Carbon.”
- Howard, Peter H and Thomas Sterner (2014) “Loaded DICE: Refining the Meta-Analysis Approach to Calibrating Climate Damage Functions.”
- Hsiang, Solomon M (2016) “Climate Econometrics.”
- Hwang, IC, Frederic Reynes, and Richard Tol (2014) “The Effect of Learning on Climate Policy Under Fat-Tailed Uncertainty,” No. 53681.
- Jensen, Sverre and CP Traeger (2013) “Optimally Climate Sensitive Policy Under Uncertainty and Learning.”

- Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero (2014) “Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain,” *Journal of Economic Dynamics and Control*, Vol. 44, pp. 92–123.
- Karp, Larry and Christian Traeger (2013) *Dynamic Methods in Environmental and Resource Economics*.
- Kelly, David L. and Charles D. Kolstad (1999) “Bayesian Learning, Growth, and Pollution,” *Journal of Economic Dynamics and Control*, Vol. 23, pp. 491–518.
- Kelly, David L and Zhuo Tan (2015) “Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails,” *Journal of Environmental Economics and Management*, Vol. 72, pp. 98–122.
- Kimball, Miles S. (1990) “Precautionary Saving in the Small and in the Large,” *Econometrica*, Vol. 58, No. 1, pp. 53–73.
- Kopp, Robert E, Alexander Golub, Nathaniel O Keohane, and Chikara Onda (2012) “The Influence of the Specification of Climate Change Damages on the Social Cost of Carbon,” *Economics: The Open-Access, Open-Assessment E-Journal*, Vol. 6.
- Krueger, Dirk and Felix Kubler (2004) “Computing Equilibrium in OLG Models with Stochastic Production,” *Journal of Economic Dynamics and Control*, Vol. 28, No. 7, pp. 1411–1436.
- Leach, Andrew (2007) “The Climate Change Learning Curve,” *Journal of Economic Dynamics and Control*, Vol. 31, No. 5, pp. 1728–1752.
- Leland, Hayne E. (1968) “Saving and Uncertainty: The Precautionary Demand For Saving,” *The Quarterly Journal of Economics*, Vol. 82, No. 3, pp. 465–473.
- Lemoine, Derek and Haewon C McJeon (2013) “Trapped Between Two Tails: Trading Off Scientific Uncertainties via Climate Targets,” *Environmental Research Letters*, Vol. 8, No. 3.
- Lemoine, Derek and Ivan Rudik (2017) “Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point,” *Annual Review of Resource Economics*, Vol. 9.

- Lemoine, Derek and Christian Traeger (2014) “Watch Your Step: Optimal Policy in a Tipping Climate,” *American Economic Journal: Economic Policy*, Vol. 6, No. 1, pp. 137–166.
- (2016a) “Ambiguous Tipping Points,” *Journal of Economic Behavior & Organization*.
- (2016b) “Economics of Tipping the Climate Dominoes,” *Nature Climate Change*, Vol. 6, No. 5, pp. 514–519.
- Lenton, Timothy M. and Juan-Carlos Ciscar (2012) “Integrating Tipping Points into Climate Impact Assessments,” *Climatic Change*, Vol. 117, No. 3, pp. 585–597.
- Lontzek, Thomas S., Yongyang Cai, Kenneth L. Judd, and Timothy M. Lenton (2015) “Stochastic Integrated Assessment of Climate Tipping Points Indicates the Need for Strict Climate Policy,” *Nature Climate Change*, Vol. 5, No. 5, pp. 441–444.
- Mirrlees, James A and Nicholas H Stern (1972) “Fairly good plans,” *Journal of Economic Theory*, Vol. 4, No. 2, pp. 268–288.
- Moore, Frances C. and Delavane B. Diaz (2015) “Temperature Impacts on Economic Growth Warrant Stringent Mitigation Policy,” *Nature Climate Change*, No. January, pp. 1–5.
- Nordhaus, W and Paul Sztorc (2013) “DICE 2013R: Introduction and User’s Manual.”
- Nordhaus, William D. (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven: Yale University Press.
- Pindyck, Robert S (2013) “Climate Change Policy: What Do the Models Tell Us?,” *Journal of Economic Literature*, Vol. 51, No. 3, pp. 860–872.
- Roseta-Palma, C and A Xepapadeas (2004) “Robust Control in Water Management,” *Journal of Risk and Uncertainty*, Vol. 29, No. 1, pp. 21–34.
- Smolyak, S. (1963) “Quadrature and Interpolation Formulas for Tensor Products of Certain Classes of Functions,” *Soviet Mathematics, Doklady*, Vol. 4, pp. 240–243.
- Stanton, Elizabeth A., Frank Ackerman, and Sivan Kartha (2009) “Inside the Integrated Assessment Models: Four Issues in Climate Economics,” *Climate and Development*, Vol. 1, No. 2, p. 166.

- Stern, Nicholas (2006) “The Stern Review on the Economic Effects of Climate Change.”
- Sterner, T. and U. M. Persson (2008) “An Even Sterner Review: Introducing Relative Prices into the Discounting Debate,” *Review of Environmental Economics and Policy*, Vol. 2, No. 1, pp. 61–76.
- Temzelides, Ted, Xin Li, and Borgahn Narajabad (2014) “Robust Dynamic Optimal Taxation and Environmental Externalities.”
- Tol, Richard S. J. (2009) “The Economic Effects of Climate Change,” *The Journal of Economic Perspectives*, Vol. 23, No. 2, pp. 29–51.
- Traeger, Christian P. (2014) “A 4-States DICE: Quantitatively Addressing Uncertainty Effects in Climate Change,” *Environmental and Resource Economics*, Vol. 59, No. 1, pp. 1–37.
- Weitzman, Martin L. (2012) “GHG Targets as Insurance Against Catastrophic Climate Damages,” *Journal of Public Economic Theory*, Vol. 14, No. 2, pp. 221–244.
- Winschel, Viktor and Markus Kratzig (2010) “Solving, Estimating, and Selecting Nonlinear Dynamic Models Without the Curse of Dimensionality,” *Econometrica*, Vol. 78, No. 2, pp. 803–821.

Online Appendix

The online appendix gives full details on the dynamic stochastic version of the DICE model, provides additional details on Robust Control, and finally, shows the results of accuracy and sensitivity checks.

A The Full Stochastic DICE Model

The DICE model is a Ramsey-Cass-Koopmans growth model coupled to a climate system. The model is governed by a representative policymaker whose objective is to maximize her expected discounted welfare. Each period in DICE lasts one decade and the model begins in 2005.

In each period t , the policymaker has an endowment of capital K_t , labor L_t , and technology A_t . These are combined in a Cobb-Douglas production function to produce gross output,

$$Y_t^g = A_t L_t^{1-\kappa} K_t^\kappa.$$

Warming of the Earth's surface causes damage to output, resulting in net output after damages,

$$Y_t^n = Y_t^g \mathcal{L}_t,$$

where \mathcal{L}_t is the fraction of output remaining after damages,

$$\mathcal{L}_t = \frac{1}{1 + d_1 [T_t^s]^{d_2} w_t}.$$

The policymaker has three ways to use her net output each period. First, she can use it for consumption C_t to increase utility,

$$U(C_t, L_t) = L_t \frac{(C_t/L_t)^{1-\eta}}{(1-\eta)}, \quad \eta \neq 1.$$

Second, she can use it to abate some fraction of emissions from factor production, α_t . The residual is left for investment into increasing the future capital stock, which depreciates at

an annual rate of δ_k . Net emissions after abatement is given by,

$$e_t = 10 [\sigma_t(1 - \alpha_t)Y_t^g + B_t].$$

B_t is emissions from exogenous land use change, and σ_t is the emissions intensity of output. Emissions enter the atmospheric stock of CO₂, M_t^{atm} . CO₂ can then flow to the upper ocean CO₂ stock M_t^{up} , and the lower ocean CO₂ stock M_t^{lo} . Increased atmospheric CO₂ increases the level of radiative forcing,

$$F_t(M_t^{atm}) = f_{2x} \log_2(M_t^{atm}/M_{pre}) + EF_t,$$

where EF_t is exogenous forcing from non-CO₂ long-lived greenhouse gases and f_{2x} is the additional forcing from a doubling of atmospheric CO₂. Radiative forcing increases the temperature of the Earth's surface, T_t^s . Thermal energy from the surface can move into the lower ocean, T_t^o and also back from the ocean to the surface depending on which is warmer. The parameter $C1$ governs the rate of warming at the Earth's surface from additional forcing or a warmer ocean, $C3$ governs how quickly the surface loses thermal energy to the ocean, and $C4$ governs how quickly the lower ocean loses thermal energy to the surface.

Finally, each period the policymaker observes the level of surface temperature, capital, labor, technology and resulting net output. Using this information she updates her beliefs over d_1 as described in the main text.

The model's exogenous economic processes are,

$$\begin{aligned} L_t &= L_0 + (L_\infty - L_0) g_{L,t} && \text{(Labor population)} \\ g_{L,t} &= [\exp(\delta_L t) - 1] / \exp(\delta_L t) && \text{(Labor population growth rate)} \\ A_t &= A_{t-1} / (1 - g_{A,t}) && \text{(Production technology)} \\ g_{A,t} &= 10 g_{A,0} \exp(-\delta_A t) && \text{(Production technology growth rate)} \end{aligned}$$

Period $t = 0$ indicates the year 2005, and each period is ten years. The exogenous climate processes are,

$$\begin{aligned} \sigma_t &= \sigma_{t-1} / (1 - g_{\sigma,t}) && \text{(Gross emissions per unit of output)} \\ g_{\sigma,t} &= g_{\sigma,0} \exp(-\delta_\sigma t) && \text{(Growth rate of gross emissions per unit of output)} \\ \psi_t &= \frac{a_0 \sigma_t}{a_1 a_2} (a_1 - 1 + \exp(-g_\psi t)) && \text{(Abatement cost coefficient)} \end{aligned}$$

$$B_t = B_0 g_B^t \quad (\text{Non-industrial CO}_2 \text{ emissions})$$

$$EF_t = EF_0 + 0.1 (EF_1 - EF_0) \min(t, 10) \quad (\text{Exogenous forcing})$$

Table 1 reports the values of the model parameters. The calibration of the distribution over d_1 and the damage shock are described in the main text.

Table 1: The parameters of the model.

Parameter	Value	Description
A_0	0.027	Initial production technology
$g_{A,0}$	0.009	Initial growth rate of production technology
δ_A	0.001	Change in growth rate of production technology
L_0	6514	Year 2005 population (millions)
L_∞	8600	Asymptotic population (millions)
δ_L	0.35	Rate of approach to asymptotic population level
σ_0	0.13	Initial emission intensity of output (Gigatons of carbon per unit output)
$g_{\sigma,0}$	-0.073	Initial growth rate of decarbonization
δ_σ	0.003	Change in growth rate of emissions intensity
a_0	1.17	Cost of backstop technology in 2005 (\$1000 per ton of carbon)
a_1	2	Ratio of initial backstop technology cost to final backstop technology cost
a_2	2.8	Abatement cost function exponent
g_Ψ	0.05	Growth rate of backstop technology cost
B_0	1.1	Initial non-industrial CO ₂ emissions (Gigatons of carbon)
g_B	0.9	Growth rate of non-industrial emissions
d_2	2	Damage exponent
EF_0	-0.06	Year 2005 exogenous forcing (W/m ²)
EF_{100}	0.30	Year 2105 exogenous forcing (W/m ²)
κ	0.3	Capital elasticity in production
δ_κ	0.1	Annual capital depreciation rate
M_{pre}	596.4	Pre-industrial atmospheric CO ₂ (Gigatons of carbon)
β	1/1.015 ¹⁰	Discount factor
η	2	1/EIS, and RRA
ϕ_{11}	0.811	Carbon transfer coefficient for atmosphere to atmosphere
ϕ_{12}	0.189	Carbon transfer coefficient for atmosphere to upper ocean

Continued on next page

Table 1 – continued from previous page

Parameter	Value	Description
ϕ_{21}	0.097	Carbon transfer coefficient for upper ocean to atmosphere
ϕ_{22}	0.853	Carbon transfer coefficient for upper ocean to upper ocean
ϕ_{23}	0.050	Carbon transfer coefficient for upper ocean to lower ocean
ϕ_{32}	0.003	Carbon transfer coefficient for lower ocean to upper ocean
ϕ_{33}	0.997	Carbon transfer coefficient for lower ocean to lower ocean
C_1	0.22	Warming delay parameter
C_3	0.3	Parameter governing transfer of heat from ocean to surface
C_4	0.05	Parameter governing transfer of heat from surface to ocean
f_{2x}	3.8	Forcing from doubling of CO ₂ (W/m ²)
s	3	Climate sensitivity
K_0	137	Year 2005 capital (trillions of USD)
M_0^{atm}	808.9	Year 2005 atmospheric CO ₂ (Gigatons of carbon)
M_0^{up}	1255	Year 2005 biosphere and upper ocean CO ₂ (Gigatons of carbon)
M_0^{lo}	18365	Year 2005 lower ocean CO ₂ (Gigatons of carbon)
T_0^{atm}	0.7307	Year 2005 atmospheric temperature (Degrees Celsius)
T_0^{ocean}	0.0068	Year 2005 ocean temperature (Degrees Celsius)
μ_0	-5.9596	Year 2005 d_1 location parameter
Σ_0	0.1904	Year 2005 d_1 shape parameter
μ_w	-.0799	Damage shock location parameter
σ_w	0.1599	Damage shock shape parameter

To improve computational efficiency I make two changes of variables that do not affect the model itself but only how it is coded. I let capital and consumption be in effective technology and labor terms,

$$k_t = \frac{K_t}{A_t^{1/(1-\kappa)} L_t},$$

$$c_t = \frac{C_t}{A_t^{1/(1-\kappa)} L_t}.$$

Utility is maintained in non-effective terms so the new utility function is,

$$u(c_t; A_t, L_t) = A_t^{(1-\eta)/(1-\kappa)} L_t \frac{c_t^{1-\eta}}{1-\eta}.$$

The policymaker's problem is then,¹

$$V_t(k_t, T_t^s, T_t^o, M_t^{atm}, M_t^{up}, M_t^{lo}, \mathcal{L}_t, \mu_t, \Sigma_t) = \max_{c_t, \alpha_t} \left\{ u(c_t; A_t, L_t) + \beta E [V_{t+1}(k_{t+1}, T_{t+1}^s, T_{t+1}^o, M_{t+1}^{atm}, M_{t+1}^{up}, M_{t+1}^{lo}, \mathcal{L}_{t+1}, \mu_{t+1}, \Sigma_{t+1})] \right\}$$

subject to transitions:

$$k_{t+1} = \frac{1}{A_{t+1}^{1/(1-\kappa)} L_{t+1}} \left[(1 - \delta_k)^{10} A_t^{1/(1-\kappa)} L_t k_t + 10 \left((1 - \psi_t \alpha_t^{a_2}) Y_t^n - A_t^{1/(1-\kappa)} L_t c_t \right) \right],$$

$$\begin{bmatrix} M_{t+1}^{atm} \\ M_{t+1}^{up} \\ M_{t+1}^{lo} \end{bmatrix} = \begin{bmatrix} \phi_{11} & \phi_{21} & 0 \\ \phi_{12} & \phi_{22} & \phi_{32} \\ 0 & \phi_{23} & \phi_{33} \end{bmatrix} \begin{bmatrix} M_t^{atm} \\ M_t^{up} \\ M_t^{lo} \end{bmatrix} + \begin{bmatrix} e_t \\ 0 \\ 0 \end{bmatrix},$$

$$T_{t+1}^s = T_t^s + C_1 \left[F_{t+1}(M_{t+1}^{atm}) - \frac{f_{2x}}{s} T_t^s + C_3 (T_t^o - T_t^s) \right],$$

$$T_{t+1}^o = C_4 T_t^s + (1 - C_4) T_t^o,$$

$$\mathcal{L}_{t+1} = \frac{1}{1 + d_1 [T_{t+1}^s]^{d_2} w_{t+1}},$$

$$\mu_{t+1} = \frac{\Sigma_t \left(\log \left(\frac{Y_t^g}{Y_t^n} - 1 \right) - \log ([T_t^s]^2) - \mu_\omega \right) + \sigma_w^2 \mu_t}{\Sigma_t + \sigma_w^2},$$

$$\Sigma_{t+1} = \frac{\Sigma_t \sigma_w^2}{\Sigma_t + \sigma_w^2}.$$

Finally I constrain abatement to be less than 100 percent and I impose the resource constraint,

$$\alpha_t \leq 1,$$

$$A_t^{1/(1-\kappa)} L_t c_t + (\psi_t \alpha_t^{a_2}) Y_t^n \leq Y_t^n.$$

If the policymaker does not learn, then $\mu_{t+1} = \mu_t$ and $\Sigma_{t+1} = \Sigma_t$.

¹ Σ_t does not technically enter the state space for collocation but I display it here for clarity.

A.1 Model Solution Method

The model is solved using value function iteration on a finite horizon. The collocation grid and polynomial interpolant are built using the Smolyak sparse grid method (Smolyak, 1963; Judd et al., 2014). The Smolyak method is optimal in the sense that, for a given number of grid points, it yields the collocation grid and set of basis functions that minimize the \mathcal{L}_2 and \mathcal{L}_∞ norms of the approximation error for the class of functions with bounded second-order mixed derivatives (Brumm and Scheidegger, 2015). Here I use an approximation level of $\mu = 3$ for the Smolyak algorithm which implies nine unique collocation points on each dimension, although the full tensor product is not used to construct the collocation grid. The terminal year is 2605. The terminal continuation value function corresponding to 2615 has the policymaker not learning while holding her initial 2005 beliefs, and has all exogenous processes held constant at their 2615 levels. Changing the terminal value function to one where the policymaker does not expect damage stochasticity and believes the damage function to be exactly equal to that in the conventional DICE model does not significantly alter the results. The relative errors between the optimal trajectories are generally on the order of 10^{-4} or smaller, but peak near 10^{-3} after 200 years as the policymaker gets closer to the terminal value function.

Expectations over future states are taken using Gauss-Hermite quadrature with seven quadrature points for both the prior and the damage shock for a total of 49 quadrature points. Increasing the number of quadrature points has virtually zero effect on the results. Also note that the support of the location parameter for a lognormal distribution is the entire real line. So, depending on the draw of d_1 and the sequence of shocks, it can take on any value in $(-\infty, \infty)$. I select a domain where in practice, virtually all of the simulation runs stay within the domain. Expanding the size of the domain used in the main text by 50 percent leads to a change in the optimal carbon tax of the learning framework of only 5×10^{-3} on average. Alternatively, using the learning domain for μ_t , or one that is 90 percent smaller, only results in a relative difference of 4×10^{-6} for the uncertainty framework's optimal carbon tax. Without truncating the distribution, I cannot guarantee that the location parameter will always stay within the domain. If it does jump outside the collocation domain during one of the simulation runs I throw out the run and perform a replacement run, however this is a rare occurrence.

B Robust Control

The policymaker faces objective randomness in how damages realize over time, but she also faces subjective uncertainty due to her limited knowledge of how temperature actually manifests as damages. Her attitudes towards the first, objective source of uncertainty, may be different than her attitudes towards the second, subjective source of uncertainty, which is generally called ambiguity. The disconnect in these two forms of uncertainty is introduced into the axiomatization of preferences as a weakening of the Axiom of Independence.

Robust control puts ambiguity into context as a concern about model misspecification. The ways her model is misspecified are unknown and unlikely to be learned in a reasonable time frame, so she is unable to apply probability distributions over possible true models and stay in the usual subjective expected utility setting. For example, assuming the damage function is truly a polynomial, she may not know the exact degree of polynomial. To capture attitudes towards ambiguity stemming from model misspecification, robust control introduces a new set of preferences for the policymaker: multiplier preferences (Strzalecki, 2011). Multiplier preferences are a special case of variational preferences (Maccheroni et al., 2006) where a penalty parameter θ_1 determines the policymaker's attitude towards ambiguity and also gauges how robustly she designs policy to potential misspecifications in her model. Multiplier preferences are characterized by a relative entropy function, alternatively called the Kullback-Leibler divergence, which measures the statistical distance between distributions.² The relative entropy function,

$$D_{t+1}(p_{t+1}||q_{t+1}) = \int \log \left(\frac{dp_{t+1}}{dq_{t+1}} \right) dp_{t+1},$$

enters the policymaker's Bellman equation as a minimization problem:

$$V_t(\mathbf{S}_t) = \min_{q_{t+1}} \left\{ \max_{c_t, \alpha_t} \left[\int u(c_t; A_t, L_t) dq_{t+1} + \theta_1 D_{t+1}(p_{t+1}||q_{t+1}) + \beta V_{t+1}(\mathbf{S}_{t+1}) \right] \right\}$$

Where p_{t+1} is her approximating model of the transition probabilities, and q_{t+1} is a distorted model of transition probabilities, where the size of the distortion is selected by an "evil agent" aiming to minimize the policymaker's expected welfare.³ In the robust control context, the policymaker maximizes her expected discounted welfare subject to facing the evil agent's

²In variational preferences, θ_1 and the relative entropy function are consolidated into a general function called the ambiguity index.

³The optimal policy is independent of the order of the maximization and minimization.

Table 2: Relative errors between GAMS solution and the dynamic programming solution to the DICE-2007 model on the domain used in this paper.

	Abatement Rate	Consumption	Investment	Temperature	CO ₂
Maximum Relative Error	6.4×10^{-4}	1.5×10^{-3}	9.6×10^{-3}	9.2×10^{-4}	1.1×10^{-3}
Average Relative Error	1.2×10^{-3}	8.9×10^{-4}	8.7×10^{-4}	6.2×10^{-4}	6.9×10^{-4}

distorted model. How much the model is distorted by the evil agent is modulated by the size of the penalty parameter θ_1 . Higher θ_1 's increase the evil agent's penalty for selecting a model much different than the approximating model and in the limit the evil agent's selection is the same as the approximating model since any non-zero relative entropy carries an infinite penalty. Conversely, decreasing θ_1 reduces the evil agent's cost of selecting larger distortions in terms of their relative entropy. Effectively, the penalty parameter governs the size of the cloud of models around her approximating models she deems as plausible alternatives. As shown in Hansen and Sargent (2007), the evil agent's optimal selection can be represented by a risk sensitivity operator,

$$T^1(V_{t+1}(\mathbf{S}_{t+1})|\theta_1) = -\theta_1 \log \left(E_{\mathbf{S}_{t+1}} \left[\exp \left(-\frac{V_{t+1}(\mathbf{S}_{t+1})}{\theta_1} \right) \right] \right).$$

If $\theta_1 = \infty$ the risk sensitivity operator reduces to the standard expected continuation value, but if $\theta_1 < \infty$, then the distribution underlying the future value is twisted by the evil agent's distortion. The risk sensitivity operator is decreasing in θ_1 , so smaller selections result in lower (perceived) future welfare. This induces the policymaker to design policy as if she is facing worse futures and helps her guard against unknown misspecifications.

C Error Analysis

I examine the error in the value function approximation by comparing the solution to the standard DICE-2007 model using the dynamic programming model and compare it against the solution reported by Nordhaus which was computed in GAMS. I report the dynamic programming solution when using the same collocation domain as the main results. To make an apples-to-apples comparison, I change the GAMS code so that the model is Markov. The standard DICE-2007 model has time t forcing as a function of the average of times t and $t + 1$ atmospheric CO₂. I alter the code so that it is instead just a function of time t

Table 3: Average and maximum first-order condition errors of 200 year simulations for each framework. The average and maximum trajectory errors are averaged over draws of d_1 , damage shock vectors, and an equal weighting across all four damage functional forms in the main text.

	Uncertainty	Learning	Robust	Robust Control and Learning
Maximum Error	6.6×10^{-4}	6.6×10^{-4}	6.8×10^{-4}	6.8×10^{-4}
Average Error	1.0×10^{-4}	1.0×10^{-4}	1.2×10^{-4}	1.0×10^{-4}

atmospheric CO₂. Relative errors between the two approaches for policy trajectories, and surface temperature and atmospheric CO₂ trajectories, are reported in Table 2.

The optimal abatement, consumption and investment policies are all very close to the GAMS solution, with differences generally on the order of tenths to hundredths of a percent. Given the high accuracy of the optimal policies, the trajectories for surface temperature and atmospheric CO₂ are also accurate and generally have errors on the order of hundredths of a percent compared to the GAMS solution.

Next I investigate internal consistency of the model by calculating the errors in simulated policy trajectories. Table 3 displays the maximum and average errors of the simulated trajectories' first-order conditions (FOCs). Time t errors in a given simulation run are defined as the infinity norm of the absolute error in the first-order conditions. Along each 200 year trajectory, the maximum and average error is then taken. These are then averaged over 20000 simulations: 5000 for each damage functional form examined in the main text.

Average errors for all frameworks are on the order of 10^{-4} , while the maximum error over all the simulations for a given framework is six times larger but still small.

D Sensitivity Analysis of the Prior and Shock Distributions

Figure 1 depicts the mean optimal carbon taxes for the uncertainty and learning frameworks when the damage function has been correctly specified by the policymaker, but when facing different variances of the prior and shock to examine the sensitivity of policies to these choices. The means of the two distributions are still held to the values in the main text. The top row shows the taxes for the uncertainty framework and the bottom row shows the taxes for the learning framework. The left column shows taxes when the shock shape parameter

has been altered by ± 50 percent, which correspond to approximately a doubling or halving of the variance. The right column shows trajectories when the prior variance is varied by ± 50 percent.

Varying the prior has very little impact for either the uncertainty or learning frameworks since learning occurs quickly with conjugate prior Bayesian updating, and a larger variance on the prior does not have substantial long run effects. Having a higher variance prior results in a minutely smaller carbon tax since higher variance in damages results in higher expected output and welfare. Altering the shock variance for the uncertainty framework again has a very small effect that works in the same way as altering the prior variance.

Altering the shock variance does have a slightly larger effect when learning, but the effect is negligible until the end of the first century where it about $\$1/tCO_2$. Altering the shock distribution also does not affect the qualitative results regarding the optimal tax decomposition and ex-post welfare.

References

- Brumm, Johannes and Simon Scheidegger (2015) “Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models.”
- Hansen, Lars Peter and Thomas J. Sargent (2007) “Recursive Robust Estimation and Control Without Commitment,” *Journal of Economic Theory*, Vol. 136, No. 1, pp. 1–27.
- Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero (2014) “Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain,” *Journal of Economic Dynamics and Control*, Vol. 44, pp. 92–123.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006) “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, Vol. 74, No. 6, pp. 1447–1498.
- Smolyak, S. (1963) “Quadrature and Interpolation Formulas for Tensor Products of Certain Classes of Functions,” *Soviet Mathematics, Doklady*, Vol. 4, pp. 240–243.
- Strzalecki, Tomasz (2011) “Axiomatic Foundations of Multiplier Preferences,” *Econometrica*, Vol. 79, No. 1, pp. 47–73.




Figure 1: Optimal tax trajectories for the uncertainty (top row) and learning (bottom row) frameworks when altering the damage shock shape parameter by ± 50 percent which approximately doubles and halves the variance (left column) and altering the prior variance by ± 50 percent.