Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point[*,†]

Derek Lemoine[‡] and Ivan Rudik[§]

University of Arizona Working Paper 17-04

February 2017

Uncertainty is critical to questions about climate change policy. Recently developed recursive integrated assessment models have become the primary tools for studying and quantifying the policy implications of uncertainty. We decompose the channels through which uncertainty affects policy and quantify them in a recursive extension of a benchmark integrated assessment model. The first wave of recursive models has made valuable, pioneering efforts at analyzing disparate sources of uncertainty. We argue that frontier numerical methods will enable the next generation of recursive models to better capture the information structure of climate change and to thereby ask new types of questions about climate change policy.

JEL: C61, D91, H23, Q54, Q58

Keywords: dynamic programming, insurance, learning, precaution, prudence, uncertainty

[*]Prepared for the *Annual Review of Resource Economics*. When citing this paper, please use the following: Lemoine, D. and I. Rudik. 2017. Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. Annu. Rev. Res. Econ. 9: Submitted. Doi: 10.1146/annurev-resource-100516-053516.

[†]We thank Paul Fackler, Jesús Fernández-Villaverde, Larry Karp, Stephen Newbold, Stan Reynolds, and Christian Traeger for helpful comments.

[‡]Department of Economics, University of Arizona, dlemoine@email.arizona.edu

[§]Department of Economics and Center for Agricultural and Rural Development, Iowa State University, irudik@iastate.edu

Economists view climate change as resulting from a massive market failure: today's greenhouse gas emissions generate global warming that will affect people and ecosystems for many generations to come, yet those emissions often carry a market price of zero. Economists therefore view the primary goal of climate change policy as pricing greenhouse gas emissions to account for the costs of the climate change they generate. In a valiant attempt to quantify these costs, economists have developed integrated assessment models (IAMs) that couple climatic and economic modules.[1] These models determine the optimal trajectory and price of greenhouse gas emissions by trading off the benefits of allowing more emissions today against the costs of future climate change. However, accurately quantifying the costs of future climate change is an impossible task: these costs depend, among much else, on the uncertain unfolding of climate change, on the uncertain consequences of climate change for the economy and for wellbeing, and on the uncertain evolution of technology and the economy. Conventional IAMs can provide insight into which parameters are likely to be important for the optimal emission price, but their deterministic structure can only take one so far in a world of uncertainty.

A newer generation of IAMs incorporates uncertainty into the optimal emission price. These "recursive" IAMs solve a dynamic programming version of a standard IAM. Their modeled policymaker is therefore cognizant of uncertainty and also of the possibility of learning through future observations of the climate and the economy. In principle, the recursive approach to IAMs can incorporate all of the uncertainties that bedevil the application of standard IAMs to policy, although in practice many of these uncertainties are of such a "deep" nature that they are difficult to formalize.

The present review outlines the recursive approach to climate change policy, demonstrates how to use recursive modeling to generate deeper theoretical insight, and summarizes the main conclusions of the recent literature. We judge this literature to be at an inflection point. The literature has made enormous strides from a standing start. However, the marginal benefit of yet another model incorporating one more standalone source of uncertainty is now low. To date, most recursive models have been limited by the use of numerical methods that are advanced relative to much work in economics but nonetheless are not well-adapted to the high-dimensional state spaces that characterize IAMs.[2] We propose that adopting methods even closer to the frontier of computational economics will allow recursive IAMs to be more detailed and representative of reality and thereby expand the types of questions that they can explore. To this end, we describe the standard numerical approach to solving a recursive IAM and provide a guide to some promising numerical methods that are not yet common in economics research. We hope that a second wave of recursive IAMs will use these methods to ask new types of questions.

[1]IAMs' estimates of the social cost of greenhouse gas emissions have formed the backbone of U.S. policy (Greenstone et al., 2011). See Nordhaus (2013) and Kelly and Kolstad (1999a) for overviews of climate-economy integrated assessment.

[2]IAMs' climate modules require several state variables on top of those required by their economic modules.

In order to motivate the recursive approach to policy under uncertainty, the next section formally describes the shortcomings of what some may see as the main alternative: the Monte Carlo approach to policy under uncertainty. Section 2 then formalizes the recursive approach. Section 3 decomposes the channels through which uncertainty about future warming affects optimal policy, and Section 4 quantifies these channels in a recursive extension of the benchmark DICE IAM (which we will make publicly available at `https://github.com/irudik/dynamic-stochastic-dice`). We hope that these two sections provide future literature with a guide to using recursive IAMs to generate theoretical insight. Section 5 summarizes results from the first wave of research with recursive IAMs. Section 6 discusses numerical methods, both standard and frontier. Section 7 concludes with suggested directions for future research. The appendix extends the theoretical analysis to the case of Epstein-Zin-Weil preferences, provides the full equations for our recursive IAM, and describes best practices for validating solutions to recursive models.

1 Shortcomings of the Monte Carlo Approach to Policy Under Uncertainty

For many years, economists numerically analyzed the implications of uncertainty for climate policy by undertaking Monte Carlo analyses of integrated assessment models. A newer, growing literature analyzes uncertainty by constructing recursive, dynamic programming versions of integrated assessment models. We begin by describing the simplest Monte Carlo approach and its shortcomings for analyzing most questions of interest about the policy implications of uncertainty. We describe the dynamic programming approach in the next section.

Consider a stylized, simplified integrated assessment model. The later numerical analysis will use a full IAM, with more states and controls. For now, the policymaker controls emissions e_t, trading off current utility from the consumption enabled by additional emissions against the welfare loss from triggering additional warming in the future. The policymaker's time t per-period utility is $u_t(e_t; T_t)$, with utility increasing and concave in emissions ($\partial u_t/\partial e_t > 0, \partial^2 u_t/\partial e_t^2 < 0$) and decreasing in temperature ($\partial u_t/\partial T_t < 0$). We allow utility to depend on time to reflect the potential for exogenous growth in population, capital stocks, or technology. Temperature evolves as $T_{t+1} = f(T_t, e_t, \epsilon_{t+1}; s)$. Increasing time t emissions increases time $t+1$ temperature ($\partial f/\partial e_t \geq 0$), and increasing time t emissions can also increase temperature after time $t+1$ via the dependence of f on T_t. The parameter s controls the climate's responsiveness to emissions (commonly referred to as "climate sensitivity"), with $\partial^2 f/\partial e_t \partial s, \partial^2 f/\partial T_t \partial s \geq 0$. The integrated assessment modeler does not know the true value of s and is interested in the implications of uncertainty about s for welfare and for emission policy. The integrated assessment modeler estimates the mean of s to be μ_0. Finally, ϵ_{t+1} is an exogenously evolving shock that is perfectly anticipated by the modeled policymaker. In later sections, we will let ϵ_{t+1} will be random and unobserved, which will prevent the policymaker from immediately learning the true value of s.

If we knew that s took on a particular value s_i with probability 1, then our policymaker could solve a deterministic problem. Specifically, the policymaker would solve:

$$W^{det}(T_0; s_i) = \max_{\{e_t\}_0^\tau} \sum_{t=0}^{\tau} \beta^t u_t(e_t; T_t), \quad \text{subject to } T_{t+1} = f(T_t, e_t, \epsilon_{t+1}; s_i),$$

where $\beta \in (0, 1)$ is the per-period discount factor, model time starts at time 0, and the policymaker's horizon extends out to $\tau > 0$, with τ potentially infinite. Solving this problem yields the policymaker's maximized welfare W^{det} as a function of initial temperature and s_i, and it also yields optimal policy $e_t^{det}(s_i)$ in every period.

Now consider a Monte Carlo analysis of uncertain s. The simplest Monte Carlo analysis of policy under uncertainty solves the deterministic problem for several different values s_i and compares the weighted average policy over these simulations to the policy resulting from fixing $s_i = \mu_0$.[3] Formally, Monte Carlo analyses define policy under uncertainty as

$$e_t^{MC} := E_0 \left[e^{det}(s_i) \right] = \int_{-\infty}^{\infty} e_t^{det}(s_i)\, p(s_i)\ \mathrm{d}s_i,$$

where E_0 refers to expectations at time 0 and $p(s_i)$ is the modeler's probability that $s = s_i$. The "Monte Carlo" label refers to techniques that approximate the expectation by taking random draws from the density $p(\cdot)$. In these analyses, the difference $e_t^{MC} - e_t^{det}(\mu_0)$ defines the effect of uncertainty on policy (e.g., Hope, 2006; Anthoff et al., 2009; Pycroft et al., 2011; Kopp et al., 2012).

Consider the information structure of the Monte Carlo analysis. Who is it that is uncertain about s? With respect to whose expectations is e_t^{MC} an average policy?

In the Monte Carlo analysis, the *modeled policymaker* is never uncertain about s. In every simulation, the policymaker knows that $s = s_i$. In particular, the policymaker is not uncertain about s when choosing emissions: emissions are chosen for a deterministic world. The emission trajectory e_t^{MC} reflects the *modeler's* uncertainty about which model to run (i.e., which value of s_i to code), but it does not reflect uncertainty within a model about which value of s_i is the correct one. This Monte Carlo analysis therefore assumes that we do

[3]Note that the simplest Monte Carlo analysis described here differs from a more sophisticated use of Monte Carlo techniques: some modelers design the policy trajectory to maximize expected welfare, with expected welfare approximated by Monte Carlo techniques (e.g., Manne and Richels, 1995; Nordhaus and Popp, 1997; Pizer, 1999; Tol, 1999; Dietz and Stern, 2015). The resulting policy is "open-loop" in that it fixes the level of time t emissions at time 0, as if the policymaker could make a one-time commitment to future emission levels. In contrast, the policies to be described in Section 2 allow time t emissions to depend on the realizations of uncertain state variables, which allows the chosen emission level to adapt to shocks and to new information.

not know the true climate sensitivity today but will know it very soon, before we get around to formulating emission control policies.[4]

Crost and Traeger (2013) show that the difference $e_t^{MC} - e_t^{det}(\mu_0)$ may not properly sign the effect of uncertainty on policy, and they also show that a Monte Carlo analysis implies that uncertainty about a damage coefficient simultaneously reduces the "optimal" abatement rate and increases the "optimal" abatement cost, two results with seemingly contradictory implications for policy. Nordhaus (1994, Chapter 8) describes similar results.

In sum, a Monte Carlo analysis can be useful for analyzing a model's sensitivity to assumptions (e.g., Hope, 2006; Nordhaus, 2008; Ackerman et al., 2010; Anthoff and Tol, 2013; Gillingham et al., 2015; Nordhaus, 2016), but averaging over policies from a deterministic model does not tell us how uncertainty affects optimal policy.

2 The Recursive Dynamic Programming Approach to Uncertainty

So how can modelers handle uncertainty? As described in footnote 3, one way modelers can correctly handle uncertainty is by constructing a policy trajectory that maximizes expected welfare, with expected welfare approximated by Monte Carlo simulations. However, this "open-loop" policy commits to all later emission levels at time 0, thereby ignoring the ability to adapt policy to shocks and to new information about unknown parameters. If a policymaker will instead have flexibility to adapt the level of emissions to unexpected changes in the world, then she can do better by developing a "closed-loop" or "feedback" policy. Such a policy defines a rule that future emissions will obey. In particular, this type of policy prescribes future emissions as a function of whatever states the climate and the economy happen to reach, so that the realized policy trajectory depends on the realized history of shocks. Recursive dynamic programming methods develop this type of responsive policy rule. We here describe the recursive dynamic programming approach to climate

[4]This Monte Carlo analysis can tell us about the expected value of perfect information (e.g., Peck and Teisberg, 1993; Nordhaus, 1994; Nordhaus and Popp, 1997): $W^{MC}(T_0) \coloneqq E_0\left[W^{det}(T_0; s_i)\right]$ gives expected welfare conditional on information arriving in the very near future, which can be contrasted with welfare under uncertainty to determine willingness to pay to acquire that information.

change.[5,6] The next section shows how to use the results of dynamic programming to study the implications of uncertainty.

Consider the same setting as in Section 1. Maximized welfare with uncertain s is

$$W^{unc}(T_0) := \max_{\{e_t(T_t)\}_0^\tau} E_0\left[\sum_{t=0}^{\tau} \beta^t u_t(e_t; T_t)\right], \quad \text{subject to } T_{t+1} = f(T_t, e_t, \epsilon_{t+1}; s).$$

Let $e_t^{unc}(T_t)$ denote optimized time t emissions as a function of time t temperature T_t. Here, we formulate policy by maximizing over the expectation of utility in every period. In the previous section, we developed policy without using expectations at all, instead taking expectations over s only along policy trajectories that had already been optimized. Here, our policy $e_t^{unc}(T_t)$ depends on both time and temperature. In the previous section, our policy e_t^{MC} depended only on time because we averaged over policies which could each perfectly

[5]A stochastic control approach can also properly address uncertainty, but it typically requires simplified types of uncertainty (such as two-point distributions) and/or shortened time horizons (such as two-period models). We here describe the benchmark approach to integrating uncertainty into a full climate-economy integrated assessment model. Recently, a literature following Golosov et al. (2014) has developed analytic IAMs that do not require advanced numerical methods. However, assumptions such as logarithmic utility combine to make the model effectively linear, which can make uncertainty uninteresting in these settings. Traeger (2015) extends the analytic setting of Golosov et al. (2014) to maintain a unitary elasticity of intertemporal substitution (i.e., to maintain a logarithmic function for aggregating per-period utility over time) while allowing for non-unitary relative risk aversion. He analyzes the welfare implications of uncertainty about the carbon cycle and about the climate system.

[6]Lemoine (2016) studies the implications of several sources of uncertainty without using an optimization-based approach, but he studies the implications for the social cost of carbon, not the optimal emission tax or the optimal quantity of emissions. The social cost of carbon is defined for a given emission trajectory, and it matches the optimal emission tax along the optimal open-loop emission trajectory.

predict future temperature.[7]

The standard way of solving the problem is to write down its recursive formulation, in the form of a Bellman equation.[8] Use $V_t(T_t, I_t)$ to denote the value of the optimal policy program at time t when temperature is T_t and information about s is encapsulated by I_t. Thus, $V_0(T_0, I_0) = W^{unc}(T_0)$. If the policymaker knows her best possible value at time $t+1$ for any given T_{t+1} and I_{t+1}, then she can select her optimal policy at time t without needing to consider policies beyond t. This insight yields the following relationship:[9]

$$V_t(T_t, I_t) = \max_{e_t} \left\{ u_t(e_t; T_t) + \beta E_t \Big[V_{t+1}(T_{t+1}, I_{t+1}) \Big] \right\}, \tag{1}$$

where T_{t+1} depends on T_t and e_t as before and where E_t denotes expectations at the time t information set I_t. Note three features of this setup, as opposed to the Monte Carlo setup. First, we here develop the optimal time t policy for any potential realization of T_t. Second, the optimal policy here accounts for the fact that future realizations $T_{t+1}, T_{t+2}, \ldots$ are uncertain. Third, we use time t information when taking expectations over the continuation value V_{t+1}. The time t policy can thus depend on information about s obtained from observing the climate prior to time t and can depend on how time t emissions might generate information about s for use in later periods.

These are clear advantages over the simplest Monte Carlo analysis described in Section 1 and also over the more sophisticated Monte Carlo analysis described at the beginning of

[7] Consider an example with $\tau = 1$. The optimal time 0 policy solves the first-order condition:

$$e_0^{unc} = \tilde{u}_0^{-1}\left(\beta \int_s \frac{\partial u_1(e_1^{unc}(T_1); T_1)}{\partial T_1} \frac{\partial T_1}{\partial e_1}\, p(s)\ \mathrm{d}s\right),$$

where $\tilde{u}_0 \coloneqq u_0'(\cdot)$. In contrast, the Monte Carlo approach from the previous section yields

$$e_0^{MC} = \int_s \tilde{u}_0^{-1}\left(\beta \frac{\partial u_1(e_1^{det}(s_i); T_1)}{\partial T_1} \frac{\partial T_1}{\partial e_1}\right) p(s)\ \mathrm{d}s.$$

These two policies are equivalent if (a) $e_1^{unc}(T_1) = e_1^{det}(s_i)$ and (b) $\tilde{u}_0^{-1}(\cdot)$ is linear. Condition (a) requires the policymaker not to adapt policy to T_1 in a way beyond how she would adapt to s_i: either e_1^{det} must be constant in s_i with $e_1^{unc}(T_1)$ constant in T_1, or each possible value of T_1 must be completely informative about s_i. Condition (a) is unlikely to hold. As for condition (b), note that the second derivative of $\tilde{u}_0^{-1}(\cdot)$ is $-\tilde{u}_0''(\cdot)/[\tilde{u}_0'(\tilde{u}_0^{-1}(\cdot))]^3$ and that $\tilde{u}_0''(\cdot) = u_0'''(\cdot)$. Therefore, condition (b) holds if and only if $u_0'''(\cdot) = 0$, so that the recursive model's agent would not undertake precautionary savings (which are never undertaken in the Monte Carlo analysis because the modeled agent is unaware of uncertainty). Most standard utility functions, including those used in integrated assessment models, have $u'''(\cdot) > 0$, so that condition (b) does not hold.

[8] Numerous textbooks derive the Bellman equation and show that its policy and value solutions match e^{unc} and W^{unc}. A standard reference is Stokey and Lucas (1989).

[9] If we used an infinite horizon setting ($\tau = \infty$) and did not allow per-period utility to vary with time, then we would have the special (and truly "recursive") case where $V_t(T_t, I_t) = V_{t+1}(T_t, I_t)$, so that we can omit the time index on the value function, writing $V(T_t, I_t)$ and $V(T_{t+1}, I_{t+1})$. We could also obtain that truly recursive representation in a nonstationary setting by tracking time in the state space.

this section. However, there is a catch: the current problem can be much more difficult to solve. Solving the Monte Carlo problem from Section 1 merely requires solving a bunch of deterministic problems, and solving the more sophisticated Monte Carlo problem described at the beginning of this section merely requires obtaining multiple samples for any candidate policy trajectory. Neither solution method is in principle much more difficult than solving the original deterministic problem. However, solving equation (1) poses a very different challenge. If we knew $V_{t+1}(\cdot, \cdot)$, the maximization could be straightforward (and almost surely simpler than maximizing over the τ-dimensional control vector in the deterministic problem): we could optimize policy at each time t separately from optimizing policy at any other time. However, we do not initially know $V_{t+1}(\cdot, \cdot)$, as it depends on maximizations at time $t+1$ and beyond. The brunt of the computational challenge rests in obtaining a good approximation of $V_{t+1}(\cdot, \cdot)$ over the relevant state space, which is just (T_t, I_t) in our stylized example but is of much higher dimension in any standard climate-economy model. Judd (1998) and Miranda and Fackler (2002) provide textbook descriptions of relevant computational methods, and recent reviews include Cai and Judd (2014), Maliar and Maliar (2014), and Fernández-Villaverde et al. (2016). We describe computational techniques in Section 6.

3 The Effect of Uncertainty on Policy

Most studies with recursive climate-economy models solve for a satisfactory approximation to each $V_t(\cdot, \cdot)$ and then use the solution to simulate trajectories for policy variables and state variables. However, the value function itself contains valuable information about the underlying drivers of policy. Teasing apart and quantifying these underlying drivers can make the results of numerical studies of climate change deeper and more generalizable. We here theoretically demonstrate the channels through which uncertainty affects policy. Only a few studies have previously used this type of information in their analysis, but we hope that this exposition will encourage future work to use the value function to learn about the underlying drivers of policy.[10] In the next section, we will quantify these drivers in an extension of the benchmark DICE integrated assessment model.

Again consider the setting described in Section 1, with ϵ_{t+1} now a stochastic shock that is independently and identically distributed over time. The value of ϵ_{t+1} will never be observed, but it will affect the observed temperature T_{t+1}. This shock in the temperature transition prevents the policymaker from learning the true value of s from a single observation of

[10] Lemoine and Traeger (2014) use the value function to tease apart the different channels through which potential tipping points affect policy. Lemoine and Traeger (2016a) quantify the channels through which aversion to ambiguity about a tipping point's threshold affects policy. Heutel et al. (2016) tease apart the channels through which tipping points and geoengineering controls interact. Rudik (2016) quantifies the different channels through which uncertainty about the damage calibration and concern for model misspecification affect policy.

temperature. This stochastic shock represents the climate's natural variability.[11]

Assume that the policymaker learns about s as a Bayesian. At time t she observes the new realization of temperature and updates her prior distribution for s to form a posterior that she uses in equation (1). Let s and ϵ_t each be normally distributed and enter the temperature transition in a linear and separable fashion. In this case (which roughly matches all of the literature to date), we have what is called a conjugate prior and thus know that the posterior distribution of s is also normal. We can write the parameters of that posterior as functions of the time t prior, of T_t, and of e_{t-1}.

Specifically, let the mean of the time t posterior for s be μ_t and the standard deviation be Σ_t. These parameters evolve according to the following closed-form equations:[12]

$$\mu_{t+1} = g(\mu_t, T_{t+1}, e_t, \epsilon_{t+1}) \quad \text{and} \quad \Sigma_{t+1} = h(\Sigma_t, e_t).$$

We here consider the standard case in which Σ_t declines deterministically.[13] The belief parameters μ_t and Σ_t do not directly affect the time t payoff, but they do affect the time t expectation operator. They are informational states that need to be tracked just like the directly payoff-relevant state T_t. The Bellman equation (1) becomes:

$$V_t(T_t, \mu_t, \Sigma_t) = \max_{e_t} \left\{ u_t(e_t; T_t) + \beta E_t \left[V_{t+1}(T_{t+1}, \mu_{t+1}, \Sigma_{t+1}) \right] \right\}, \tag{2}$$

subject to the transition equations for T_{t+1}, μ_{t+1}, and Σ_{t+1}. The policymaker forms expectations of V_{t+1} by integrating over her subjective time t beliefs about s and over the objective distribution for the noise term ϵ_{t+1}. These expectations are conditioned on the current informational states μ_t and Σ_t as well as on T_t and on the policymaker's choice of e_t.

The policymaker's optimal choice of time t emissions is governed by the following first-order condition:

$$\frac{\partial u_t(e_t; T_t)}{\partial e_t} = -\beta E_t \left[\frac{\partial V_{t+1}}{\partial T_{t+1}} \frac{\partial T_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \mu_{t+1}} \frac{\partial \mu_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \Sigma_{t+1}} \frac{\partial \Sigma_{t+1}}{\partial e_t} \right].$$

Greater marginal utility from emissions (left-hand side) corresponds to less emissions and thus to a greater optimal tax on emissions. The right-hand side of this equation gives the welfare cost of additional emissions.[14] Pass the expectation operator through to decompose

[11]This shock does not represent imperfect monitoring of the true state of the climate because the shock's realized value affects the true value of the state and thus the policymaker's utility.

[12]For specific examples, see Kelly and Tan (2015) or the appendix of this paper. See also Cyert and DeGroot (1974) for more on conjugate priors.

[13]In ongoing work with Max Rosenthal, we solve a dynamic programming setting that allows for Bayesian learning with nonconjugate priors, in which case Σ_t can be stochastic and nonmonotonic.

[14]If we explicitly wrote consumption separately from emissions, we could divide by the time t marginal utility of consumption to obtain the optimal time t tax on emissions. Also, note that we could use the

the right-hand side into its economic components:

$$\begin{aligned}\frac{\partial u_t(e_t;T_t)}{\partial e_t} =&\beta\Bigg\{\underbrace{E_t\left[\frac{-\partial V_{t+1}}{\partial T_{t+1}}\right]}_{\text{A}} E_t\left[\frac{\partial T_{t+1}}{\partial e_t}\right] + \underbrace{Cov_t\left[\frac{-\partial V_{t+1}}{\partial T_{t+1}},\ \frac{\partial T_{t+1}}{\partial e_t}\right]}_{\text{insurance}} \\ &+ \underbrace{E_t\left[\frac{-\partial V_{t+1}}{\partial \mu_{t+1}}\frac{\partial \mu_{t+1}}{\partial e_t}\right] + E_t\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\frac{\partial \Sigma_{t+1}}{\partial e_t}\right]}_{\text{active learning}}\Bigg\}. \qquad (3)\end{aligned}$$

We see four components that determine optimal time t emissions. We will soon see that the first term on the first line (labeled A) is itself composed of several components, including a precautionary abatement component and a passive learning component. The second component is insurance against the marginal effect of emissions on temperature covarying with the marginal welfare cost of temperature. The last line determines how active learning (i.e., the ability to generate information about s through the choice of emissions) affects the optimal emission policy. We proceed to analyze each of these components in turn.

3.1 Certainty-equivalence, precautionary abatement, and passive learning

Begin with the component labeled A in equation (3). This component is the expected welfare cost of additional time $t+1$ temperature. It multiplies the expected increase in time $t+1$ temperature from additional time t emissions. A second-order Taylor expansion of $E_t[-\partial V_{t+1}/\partial T_{t+1}]$ around $z := (E_t[T_{t+1}], \mu_t, \Sigma_{t+1})$ yields (noting that $E_t[\mu_{t+1}] = \mu_t$ for a Bayesian and that the first-order terms are all zero):

$$\begin{aligned}E_t\left[\frac{-\partial V_{t+1}}{\partial T_{t+1}}\right] \approx& \underbrace{\left.\frac{-\partial V_{t+1}}{\partial T_{t+1}}\right|_{(E_t[T_{t+1}],\mu_t,0)}}_{\text{certainty-equivalent}} + \underbrace{\left\{\left.\frac{-\partial V_{t+1}}{\partial T_{t+1}}\right|_z - \left.\frac{-\partial V_{t+1}}{\partial T_{t+1}}\right|_{(E_t[T_{t+1}],\mu_t,0)}\right\}}_{\text{adjustment for future uncertainty}} \\ &+ \underbrace{\left\{\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial T^3_{t+1}}\right|_z Var_t(T_{t+1}) + \left.\frac{-\partial^3 V_{t+1}}{\partial \mu_{t+1}\partial T^2_{t+1}}\right|_z Cov_t(T_{t+1},\mu_{t+1})\right\}}_{\text{precautionary abatement}} \\ &+ \underbrace{\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial \mu^2_{t+1}\partial T_{t+1}}\right|_z Var_t(\mu_{t+1})}_{\text{signal smoothing}}. \qquad (4)\end{aligned}$$

envelope theorem to substitute in for the derivatives of the value function. Repeating this procedure would yield an infinite sum on the right-hand side in terms of model fundamentals. See Jensen and Traeger (2014, 2016) for examples. We here refrain from undertaking this substitution because we will use the value function derivatives to obtain intuition that relates to familiar consumption-savings models.

The first term on the right-hand side is the certainty-equivalent tax on emissions: the welfare loss from a marginal increase in temperature is evaluated with s known to be fixed at its time t mean. This tax would be the tax in a deterministic model that happened to reach T_t at time t with $s_i = \mu_t$.[15] The term in braces on the first line adjusts the certainty-equivalent tax to use the correct continuation value. This adjustment changes the time t tax to reflect how uncertainty changes the marginal effect of temperature on future welfare. In other words, this adjustment accounts for how all of the components to be described below also affect policies and welfare after time t. This adjustment would be zero in the unrealistic case where the time t policymaker knew that s would be fixed at μ_t once time $t+1$ arrived, even though she is potentially unsure about the value of s that will govern the transition from time t to time $t+1$.

The second line on the right-hand side of equation (4) describes a precautionary abatement motive. In standard savings models, the third derivative of utility determines whether increasing uncertainty about future consumption leads an agent to save more, in which case we call the agent "prudent" (Leland, 1968; Drèze and Modigliani, 1972; Kimball, 1990). In our setting, reducing emissions (i.e., increasing "abatement") is a form of saving: it requires forgoing current utility in order to obtain additional future utility. Just as agents in standard settings undertake precautionary savings when the marginal utility of consumption is convex, here agents undertake precautionary abatement when the marginal welfare cost of temperature is convex (i.e., when $-\partial^3 V_{t+1}/\partial T_{t+1}^3 > 0$).[16,17] This precautionary abatement channel becomes stronger as the variance of temperature increases.

The covariance term in the second line of equation (4) accounts for how anticipated learning affects precautionary abatement. The sign of $-\partial^3 V_{t+1}/\partial \mu_{t+1} \partial T_{t+1}^2$ probably matches the sign of $-\partial^3 V_{t+1}/\partial T_{t+1}^3$ because the implications of raising μ_{t+1} are similar to the implications of raising T_{t+1}.[18] As shown in the transition equations for the full DICE model in the appendix, $Cov_t(T_{t+1}, \mu_{t+1}) > 0$ because high temperatures are a signal of high climate sensitivity: we tend to learn that future temperatures will drift up especially fast when we have

[15] By fixing s at its time t mean rather than its time 0 mean, we are calculating the certainty-equivalent emission tax as the tax that would be optimal conditional on reaching the time t states and then ignoring uncertainty.

[16] Repeatedly applying the envelope theorem to the Bellman equation shows that $-\partial^3 V_{t+1}/\partial T_{t+1}^3$ is closely related to the third derivative of utility. Indeed, Sibley (1975) and Carroll and Kimball (1996) show that the value function can inherit the third derivative of utility.

[17] Another way to think about the sign of $-\partial^3 V_{t+1}/\partial T_{t+1}^3$ is to consider whether the policymaker would prefer to attach a mean-zero temperature risk to a state with more severe climate change or to a state with less severe climate change (cf. Eeckhoudt et al., 1995; Crainich et al., 2013). In the latter case, we have $-\partial^3 V_{t+1}/\partial T_{t+1}^3 > 0$.

[18] We could say that $-\partial^3 V_{t+1}/\partial \mu_{t+1} \partial T_{t+1}^2 \geq 0$ when the policymaker is "cross-prudent" in temperature with respect to expected climate sensitivity: the agent prefers to attach a mean-zero temperature risk to a state with lower expected climate sensitivity than to a state with greater expected climate sensitivity. However, note that we are here considering the derivatives of the value function rather than of per-period utility. For more on cross-prudence, see Gollier (2010).

already seen temperature start to rise. Anticipating learning about the climate's sensitivity to emissions then makes future temperature appear especially variable, because the policymaker does not know what she will learn but expects to learn something that will amplify whatever the observed near-term change in temperature turns out to be. The precautionary abatement motive is therefore strengthened by this positive covariance between the posterior mean of s and temperature.

The third line of equation (4) describes another channel through which passive learning (the potentially exogenous arrival of information about s) affects the optimal emission tax. In particular, it describes how temperature affects the policymaker's ability to smooth welfare in response to whatever signal she receives about s. The curvature of V_{t+1} in μ_{t+1} captures the policymaker's ability to smooth the consequences of a higher estimate of climate sensitivity. Note that high s is bad: $-\partial V_{t+1}/\partial \mu_{t+1} > 0$. A positive value of $-\partial^2 V_{t+1}/\partial \mu_{t+1}^2$ indicates that the marginal cost of s increases in the level of s, as when damages are convex in the level of warming. The anticipated arrival of information about s makes μ_{t+1} variable from the perspective of time t: the policymaker expects to revise her beliefs when she sees the new information, but she does not know what that information will be. When $-\partial^2 V_{t+1}/\partial \mu_{t+1}^2$ $>$ $(<)$ 0, the variance $Var_t(\mu_{t+1})$ of future beliefs reduces (increases) the policymaker's expected future welfare. When $-\partial^3 V_{t+1}/\partial \mu_{t+1}^2 \partial T_{t+1} > 0$, higher temperatures hinder the policymaker's ability to smooth welfare in response to different signals about s.[19,20] In this plausible case, raising T_{t+1} (through higher e_t) increases the cost of the variance in μ_{t+1}. The policymaker then has an additional incentive to reduce emissions.[21]

3.2 Insurance

We now consider the other channels in equation (3). Begin with the insurance channel, which reduces optimal emissions if and only if $Cov_t\left[\frac{-\partial V_{t+1}}{\partial T_{t+1}}, \frac{\partial T_{t+1}}{\partial e_t}\right]$ is positive. This channel depends on whether the marginal effect of emissions on temperature tends to be high when

[19] Jones and Ostroy (1984) show that making future beliefs more variable increases the value of more flexible actions. If higher temperatures reduce the policymaker's flexibility to respond to new information about s, then anticipating learning will reduce optimal emissions. In general, there could be competing irreversibilities that make it a priori unclear whether reducing emissions increases flexibility to respond to information about different sources of uncertainty (see Fisher and Narain, 2003).

[20] We could say that $-\partial^3 V_{t+1}/\partial \mu_{t+1}^2 \partial T_{t+1} \geq 0$ when the agent is "cross-prudent" in her expectations of climate sensitivity with respect to temperature: the agent prefers to attach a mean-zero shock to expected climate sensitivity (due, for instance, to a more precise signal) to a state with lower temperature than to a state with higher temperature. However, note that we are here considering the derivatives of the value function rather than of per-period utility. For more on cross-prudence, see Gollier (2010).

[21] This signal smoothing effect is related to the wealth effect described in Gollier et al. (2000) and Gollier (2001, Chapter 25). Those analyses emphasize how increasing the informativeness of an anticipated signal induces precautionary saving but also reduces saving by increasing expected wealth, due to an enhanced ability to optimize in response to the signal. We see the precautionary effect in the $Cov_t(T_{t+1}, \mu_{t+1})$ term analyzed above.

the marginal welfare cost of temperature is high or when the marginal welfare cost of temperature is low. In the former case, additional emissions have their strongest effect on the climate when climate change matters the most for welfare. The covariance is then positive, working to reduce optimal emissions and increase the optimal emission tax. In the latter case, additional emissions have their strongest effect on the climate when climate change matters the least for welfare. The covariance is then negative, working to increase optimal emissions and reduce the optimal emission tax.[22]

What sign should we expect this covariance to have? Continue to focus on uncertainty about the climate's response to emissions.[23] If we live in a world in which the climate is very responsive to emissions (s is large), then we might expect emissions to have a strong effect on temperature ($\partial T_{t+1}/\partial e_t$ to be large) and additional warming to be especially costly ($-\partial V_{t+1}/\partial T_{t+1}$ to be large). In this case, the covariance is positive and the insurance channel increases the value of emission reductions. This is the story analyzed in previous theoretical literature (Howarth, 2003; Sandsmark and Vennemo, 2007; Becker et al., 2010; Gollier, 2012; Dietz et al., 2015; Lemoine, 2016).[24]

3.3 Active learning

Finally, consider the terms on the bottom line of equation (3). These terms are nonzero only if emission decisions affect the mean and/or the variance of future beliefs.[25]

The first term on the bottom line of equation (3) accounts for how emissions change

[22]This logic is identical to the logic driving the consumption-based capital asset pricing model (Lucas, 1978; Breeden, 1979). There, we judge stocks to be risky based not on the variance of their returns but on the covariance of their returns with marginal utility. Stocks that pay off when marginal utility is high (i.e., when consumption is low) act as valuable hedges, which reduces the expected return that investors require to hold the stocks. Stocks that pay off when marginal utility is low (i.e., when consumption is high) make consumption more volatile, which increases the expected return that investors require to hold the stocks.

[23]See Gollier (2012), Litterman (2013), Weitzman (2013), Dietz et al. (2015), and Lemoine (2016) for consideration of uncertainty about future consumption growth.

[24]The story could become more complicated once we realize that consumption damages often scale with output (Lemoine, 2016) or that large s also implies that the climate has high inertia. As described in the scientific literature (e.g., Hansen et al., 1985; Raper et al., 2002; Baker and Roe, 2009; Urban and Keller, 2009), the climate system's response time falls when s increases because oceans at first absorb much of the additional heat from an increase in carbon dioxide. This increase in inertia would complicate our intuition about the sign of the covariance, because additional inertia should affect both derivatives in the covariance.

[25]In order for emissions to affect beliefs, emissions must interact with the uncertain parameter to determine the signal that the policymaker uses to update beliefs. If time $t+1$ temperature depends on emissions and climate sensitivity in a separable fashion, then the policymaker's ability to learn about climate sensitivity is not affected by decisions about emissions. In contrast, if carbon dioxide and climate sensitivity enter the temperature transition multiplicatively, then the marginal effect of s on temperature changes with the level of emissions. In this case, greater emissions can help the policymaker to back out the true value of s.

mean beliefs. Passing the expectation operator through, we have:

$$E_t\left[\frac{-\partial V_{t+1}}{\partial \mu_{t+1}}\frac{\partial \mu_{t+1}}{\partial e_t}\right] = Cov_t\left[\frac{-\partial V_{t+1}}{\partial \mu_{t+1}},\ \frac{\partial \mu_{t+1}}{\partial e_t}\right].$$

This expression recognizes that $E_t\left[\partial \mu_{t+1}/\partial e_t\right] = 0$, because a Bayesian policymaker cannot expect to learn in a particular direction or else she should have already updated her beliefs. If states in which additional emissions lead the policymaker to substantially raise her estimate of the climate's sensitivity to emissions ($\partial \mu_{t+1}/\partial e_t$ is large) correspond to states in which these larger estimates are especially costly ($-\partial V_{t+1}/\mu_{t+1}$ is large), then the covariance is positive and the policymaker reduces optimal emissions so as to avoid learning as rapidly. This story is in fact a plausible one: we would expect additional emissions to provide better signals when s is large, and because large s corresponds to large T_{t+1}, we might expect the welfare cost of higher mean estimates to be greater when s is larger. Intuitively, generating additional information through additional emissions increases risk: time $t+1$ beliefs become more variable from the perspective of time t. If the marginal cost of mean climate sensitivity ($-\partial V_{t+1}/\partial \mu_{t+1}$) is constant, then this additional risk does not reduce expected welfare, but if the marginal cost of mean climate sensitivity is larger when the mean is higher, then this additional risk is undesirable. The policymaker then reduces emissions in order to avoid bearing the additional risk introduced by anticipated learning.

The second term on the bottom line of equation (3) captures how additional emissions generate information that allows the policymaker to develop more precise beliefs. Passing the expectation operator through and recognizing that the evolution of Σ_t is here deterministic, we have:

$$E_t\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\frac{\partial \Sigma_{t+1}}{\partial e_t}\right] = E_t\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right]\frac{\partial \Sigma_{t+1}}{\partial e_t}.$$

The term $-\partial V_{t+1}/\partial \Sigma_{t+1}$ captures the value of information. It is positive because information has positive value. The term $\partial \Sigma_{t+1}/\partial e_t$ captures how emissions affect the variance of beliefs. It is negative when additional emissions allow for faster learning, as in many integrated assessment models. In that case, this active learning channel increases optimal emissions (and reduces the optimal tax on emissions) so as to generate additional information. This is the active learning channel that informal discussions in the literature have focused on.

This second active learning channel is also the first channel that we have analyzed in which the most plausible consequence is to increase optimal emissions and thus reduce the optimal tax on emissions.[26] Therefore, when the climate's sensitivity to emissions is uncertain, we should expect this uncertainty to increase the optimal tax on emissions in the absence of learning; we should expect the introduction of passive learning (the anticipated

[26]The appendix shows that using Epstein-Zin-Weil preferences introduces a new channel that reflects preferences over the temporal resolution of uncertainty. It explains why this channel is likely to increase the optimal emission tax in standard calibrations.

and exogenous arrival of information) to further increase the optimal tax on emissions; and we should expect that the introduction of active learning (the endogenous generation of information) could reduce the optimal tax on emissions. We next quantify these channels in an extension of a benchmark integrated assessment model before reviewing the results of previous literature. Several of these channels depend on third derivatives and on cross-derivatives, so it is important to solve recursive climate-economy models with computational methods that do not a priori constrain these higher derivatives and cross-derivatives. We discuss computational approaches in Section 6.

4 Quantifying the Implications of Uncertainty for Policy

We now quantify the channels analyzed in Section 3 in a recursive extension of the benchmark DICE-2007 integrated assessment model of Nordhaus (2008). The DICE model couples a Ramsey-Cass-Koopmans growth model to calibrated modules that describe the transfer of carbon dioxide between the atmosphere and the ocean and that describe the evolution of atmospheric and oceanic temperature. Global warming reduces output directly. In each period, the policymaker chooses consumption, savings, and abatement to maximize intertemporal welfare. Savings increase future output by increasing the capital stock, and additional abatement increases future output by reducing future temperature. We extend DICE to include uncertainty about the climate's sensitivity to emissions. Following recent economic (Kelly and Tan, 2015; Hwang et al., 2017) and scientific (Roe and Baker, 2007; Roe, 2009) literature, the uncertain parameters are the feedbacks that determine the climate's sensitivity to emissions. This uncertainty can generate fat tails in the distribution of climate sensitivity. The appendix provides the complete model equations, and numerous other papers describe the DICE model in more detail. The appendix also describes how we map the analysis of Section 3 into this more complex numerical setting. Our model's code is available online at `https://github.com/irudik/dynamic-stochastic-dice`.

We find that uncertainty about the climate's sensitivity to emissions has only a small effect on the optimal emission tax. The optimal emission tax in the year 2005 would be \$7.80 per tCO_2 in the absence of uncertainty (plotted in the appendix), increasing very slightly to \$7.87 per tCO_2 in the presence of uncertainty, and increasing further to \$8.52 per tCO_2 when the policymaker anticipates learning about climate sensitivity. In the absence of uncertainty, the optimal tax would grow to \$23.25 per tCO_2 after fifty years. When climate sensitivity is uncertain (and just happens to take its mean value), the optimal tax grows to \$23.46 per tCO_2 after fifty years in the absence of learning and to \$24.57 per tCO_2 in the presence of learning.[27] Uncertainty does not have a significant effect on savings.

[27] Kelly and Tan (2015) report that uncertainty has larger consequences for emission policy. Because they

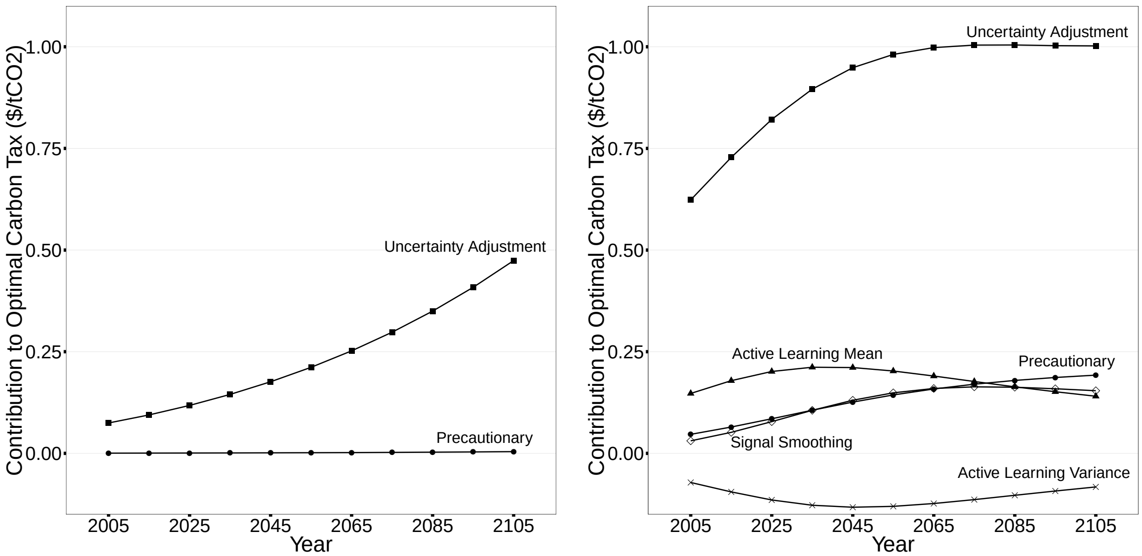

(a) Without Learning (b) With Learning

Figure 1: The channels that determine how uncertainty about the climate's sensitivity to emissions changes the optimal emission tax, for settings without learning (left) and with anticipated learning (right). All simulations fix random variables at their mean values.

We are, however, here more interested in the channels through which uncertainty matters for policy than in the total magnitude of the effect. Figure 1 plots several of the channels analyzed in Section 3. It does not include the insurance channel because it is always zero: $\partial T_{t+1}/\partial e_t$ is not uncertain because of how the DICE model formulates its temperature transition (see appendix).[28] Both panels of the figure plot the evolution of each component when climate sensitivity takes its mean value (unbeknownst to the policymaker).

The left panel plots the components in the absence of learning. In this case, the mean belief μ_t does not evolve, so the signal smoothing and active learning channels are zero. As expected, the precautionary abatement channel does increase the optimal emission tax, but only by a very small amount. The variance of temperature at time $t+1$ is always fairly small from the perspective of time t, in part because the climate system's inertia means that uncertainty about climate sensitivity takes time to manifest as temperature. The important channel is the adjustment for future uncertainty. This channel is the dominant one because

use a model that is structurally different from DICE, it is difficult to be sure where these differences stem from. Our results are similar to the small consequences of uncertainty in Jensen and Traeger (2016) and Hwang et al. (2017).

[28]Note, however, that $\partial T_{t+x}/\partial e_t$ is uncertain for $x > 1$, so the insurance channel is one of the components of the uncertainty-adjustment channel. We find that this channel is positive but very small when we consider $x = 2$. The small magnitude is likely due to the climate system's inertia.

high or low values for climate sensitivity take time to manifest as warming. Raising emissions at any given time t has a strong effect on temperature only after several periods, which shows up by changing the continuation value V_{t+1}. The uncertainty adjustment grows stronger over time as the impacts of higher emissions become more imminent. In principle, one could decompose the uncertainty adjustment into contributions from insurance and precautionary abatement motives.

The right panel plots the components when the policymaker anticipates learning. Now the mean belief μ_t evolves with observations of temperature. The adjustment for future uncertainty is still the largest single component, again reflecting the delayed consequences of climate sensitivity for temperature. The precautionary abatement channel again increases the optimal emission tax. It is larger than it was in the absence of learning, primarily because the covariance between temperature and mean beliefs is nontrivial. The precautionary abatement channel increases over the first half of the century. There are two competing effects governing how the precautionary abatement channel evolves over time. First, uncertainty about climate sensitivity declines with learning, weakening precautionary motives. Second, warming over the first few decades raises the subjective variance of future temperature and strengthens the covariance between future temperature and the mean of our future beliefs, strengthening the precautionary abatement channel. The latter effect dominates the first effect over the coming century.

We also see three channels absent from the case without learning. First, the signal smoothing channel increases the optimal emission tax because the policymaker delays emissions in order to take advantage of new information about climate sensitivity. Second, the mean-belief component of the active learning channel increases the optimal emission tax because improving the informativeness of the temperature signal (through higher emissions) creates additional risk. To our knowledge, neither of these two learning channels has been discussed in previous work with recursive models. Third, the belief-precision component of the active learning channel reduces the optimal emission tax because the policymaker can learn faster by increasing emissions, but this belief-precision component is dominated by the mean-belief component.[29] Each component of the active learning channel strengthens over the first several decades as the variance of temperature and thus the variance of the mean of future beliefs each increase, before beginning to dissipate as the policymaker hones in on the true climate sensitivity. The plotted signs for the signal smoothing channel and for

[29] In the DICE model, both parts of the active learning channel are implicitly captured by all the other channels when valuing the impact of time t emissions on the time $t+1$ expected continuation value. The appendix shows how the time $t+1$ variance and mean of the uncertain parameter are functions of time t temperature, which is a function of time $t-1$ emissions, so that the active learning channel is only explicitly present when looking two periods ahead. To assess the active learning channel, we plot $-\beta\, E_t\left[\frac{\partial V_{t+2}}{\partial \mu_{t+2}}\,\frac{\partial \mu_{t+2}}{\partial T_{t+1}}\,\frac{\partial T_{t+1}}{\partial e_t}\right]$ and $-\beta\, E_t\left[\frac{\partial V_{t+2}}{\partial \Sigma_{t+2}}\,\frac{\partial \Sigma_{t+2}}{\partial T_{t+1}}\,\frac{\partial T_{t+1}}{\partial e_t}\right]$. Note that this channel is not completely disentangled from the other, one-step-ahead channels. The optimal time t emission tax is the sum of the other plotted components (i.e., not including the active learning channel) and the certainty-equivalent tax.

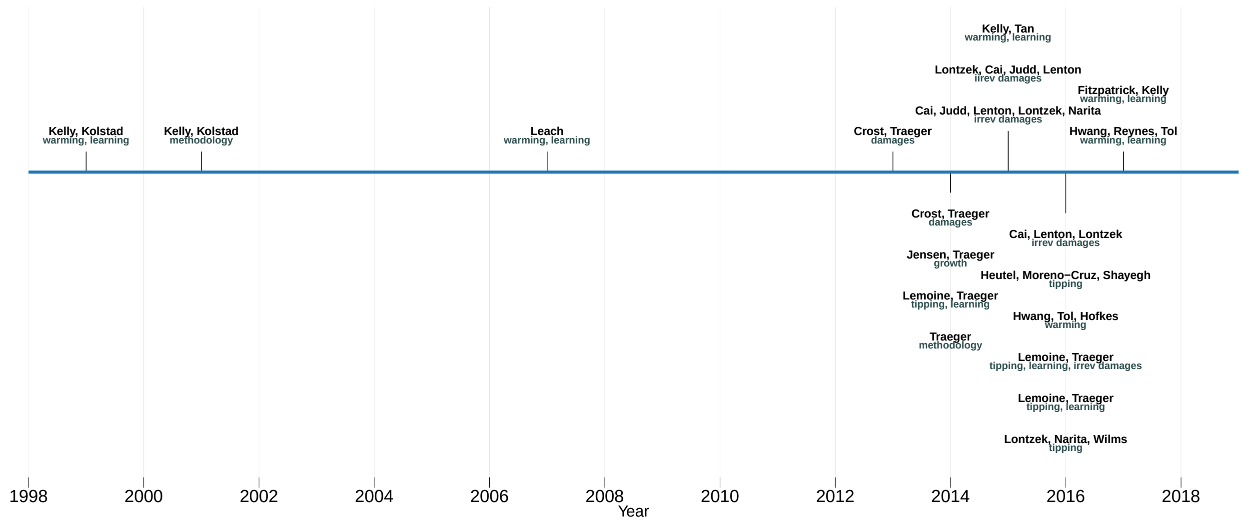

Figure 2: Publications that use recursive climate-economy integrated assessment models.

the active learning channel's components were considered to be the most likely signs in the discussion in Section 3.

We have seen that passive learning does in fact strengthen the precautionary abatement channel, that the signal smoothing channel is in fact positive, and that the active learning channel nets out to a positive value. Thus, for a given trajectory of information, the introduction of passive learning works to raise the optimal emission tax and making learning active then raises the optimal emission tax by a bit more.

5 The First Wave of Research with Recursive Models

We have thus far demonstrated how to use the value function to extract information about the channels through which uncertainty affects policy. We now review the main takeaways from previous work before discussing computational methods and suggesting directions for future work.

Figure 2 provides a timeline of published work that uses recursive IAMs. Kelly and Kolstad (1999b) were far ahead of the rest of the literature. Only two other papers were published using recursive techniques before 2013: a methodological paper by the same authors (Kelly and Kolstad, 2001), and an extension to a more realistic model of learning (Leach, 2007). Christian Traeger's research led the more recent development of recursive IAMs. He published five papers (with various coauthors) in this decade before anybody else had published one. Now recursive IAMs are hot. We count sixteen published papers and seven unpublished working papers since 2013, with only three published papers prior to 2013.

We now summarize what we see as the main lessons of this work. We focus on qualitative takeaways because model structures differ across papers and because these papers do not all report comparable experiments.

1. Uncertainty about the climate's sensitivity to emissions strongly reduces optimal emissions in one setting (Kelly and Tan, 2015) but reduces optimal emissions by only a small amount in other settings (Jensen and Traeger, 2016; Hwang et al., 2017). Anticipating the ability to learn about the climate's sensitivity to emissions strongly reduces the effect of uncertainty on emission policy in one setting (Kelly and Tan, 2015), slightly reduces the effect of uncertainty on emission policy in another setting (Hwang et al., 2017), and has basically no effect on emission policy in a third setting (Jensen and Traeger, 2016). Uncertainty about warming also increases the cost of agreeing to inflexible limits on total warming rather than using new information to control policy optimally (Fitzpatrick and Kelly, 2017). In the present paper, we have seen that uncertainty about climate sensitivity has only a small effect on emission policy and that anticipating learning can strengthen emission policy. Because all of these papers use slightly different settings and solution methods, it is not clear what drives their different results.

2. Other sources of uncertainty have been studied in only a single setting thus far. First, Crost and Traeger (2014) study uncertainty about the scaling coefficient and the exponent in the DICE damage function, which reduces output multiplicatively. They find that uncertainty about the scaling coefficient increases optimal emissions by a small amount because it increases expected output in that particular damage representation, whereas uncertainty about the damage exponent reduces optimal emissions (and to an overall greater degree) because it emphasizes high-damage realizations and thus reduces expected output. Second, Jensen and Traeger (2014) show that uncertainty about future economic growth reduces optimal emissions slightly if the shocks to growth are not persistent and reduces emissions to a greater degree if the shocks to growth are persistent. Under Epstein-Zin-Weil preferences (see below), uncertainty about economic growth increases optimal emissions.[30]

3. Potential tipping points increase the optimal emission tax. Tipping points have been represented as triggering discontinuous changes in the parameters governing either the

[30] Lemoine (2016) shows that the literature's tendency to isolate a single source of uncertainty at a time may severely underestimate the implications of uncertainty for policy. He directly compares the implications of uncertainty about warming, uncertainty about the degree to which warming reduces the growth rate of consumption, uncertainty about the future variability of the weather, and uncertainty about future consumption growth for the social cost of carbon, rather than for the optimal emission tax. He shows that uncertainty increases the social cost of carbon by a large amount under standard preferences, driven by uncertainty about the damages from additional warming as well as by the interaction between uncertainty about damages and uncertainty about warming.

carbon cycle or the climate system. Triggering a tipping point therefore causes discontinuous changes in state variables' transition equations but not in the state variables themselves. Prior to tipping, the magnitude of the increase in the optimal emission tax depends on the physical consequences of tipping (Lemoine and Traeger, 2014), on the potential for interactions among tipping points (Lemoine and Traeger, 2016b), on assumptions about the policymaker's ability to learn about tipping points prior to triggering them (Lemoine and Traeger, 2014), and on assumptions about the policymaker's ability to respond to tipping points after having triggered them (Lemoine and Traeger, 2016b). The policy implications of potential tipping points also depend on whether deploying "geoengineering" technologies can mitigate the consequences of tipping (Heutel et al., 2016).

4. Allowing warming to raise the chance of irreversible reductions in economic output (Cai et al., 2013; Lontzek et al., 2015; Cai et al., 2016) or in environmental quality (Cai et al., 2015) can strongly reduce optimal emissions. This effect is especially large under common calibrations of recursive utility (Cai et al., 2013, 2016).

5. Using common calibrations of recursive utility strongly reduces optimal emissions. Recursive utility, often implemented as Epstein-Zin-Weil (Epstein and Zin, 1989; Weil, 1989) preferences, separates risk aversion from the elasticity of intertemporal substitution, whereas the standard expected utility representation forces one to be the reciprocal of the other. Common calibrations set the elasticity of intertemporal substitution to a value strictly greater than 1 and set relative risk aversion to an even greater value. Much work has shown that increasing the elasticity of intertemporal substitution from standard expected utility calibrations strongly increases the optimal emission tax by reducing the consumption discount rate, and the high level of risk aversion (along with the implied preference for an early resolution of uncertainty) can also make policy especially sensitive to uncertainty (Cai et al., 2013; Crost and Traeger, 2014; Jensen and Traeger, 2014; Cai et al., 2016). The appendix extends the analysis of Section 3 to the case of recursive utility.

6. Other attitudes towards uncertainty matter less than does the adoption of recursive utility, probably because they do not affect the consumption discount rate as directly. In particular, a preference for robustness to alternative damage functions (Rudik, 2016) and aversion to ambiguity about a potential tipping point (Lemoine and Traeger, 2016a) each affect policy to a smaller degree than does the adoption of recursive utility.

6 Computational Discussion

The first wave of recursive IAMs explored the implications of different types of uncertainty for optimal emissions. We here describe the computational techniques used in this literature and point to improvements that could enable a second wave of research to ask questions that require more complicated models.

The standard solution technique used in the literature is value function iteration. Begin by supposing we have an infinite-horizon problem with stationary per-period utility. Using the setting of Section 3 and ignoring the potential for learning (so that we can simplify the exposition by ignoring the informational states in I_t), the Bellman equation (1) becomes

$$V(T_t) = \max_{e_t} \left\{ u(e_t; T_t) + E_t \Big[\beta\, V(T_{t+1}) \Big] \right\}, \quad \text{subject to } T_{t+1} = f(T_t, e_t, \epsilon_{t+1}; s). \tag{5}$$

A unique solution $V(T_t)$ to equation (5) exists if u is real-valued, continuous, and bounded; $\beta \in (0, 1)$; and the feasible set of states for the next period is compact (Judd, 1998). Alternative sets of assumptions that guarantee a unique solution can be found in Stokey and Lucas (1989). We can ensure that the solution is approached in the limit as $j \to \infty$ by iterating on V as follows:[31]

$$V_{j+1}(T_t) = \max_{e_t} \left\{ u(e_t; T_t) + E_t \Big[\beta\, V_j(T_{t+1}) \Big] \right\}.$$

If we have some arbitrary value function V_j, we can insert it into the right-hand side of the Bellman equation and undertake the maximization to recover a new value function V_{j+1}. If we repeat this process, we can eventually converge to the true value function for any initial V_0 (Judd, 1998). In practice, however, this algorithm may not work perfectly due to numerical error in the maximization step or if the problem is ill-conditioned.

The first critical question is: how do we form a value function V_j? The conventional approach is to approximate V with a set of basis functions $\phi(T)$:[32]

$$\tilde{V}_j(T) = \sum_{i=1}^{N} c_j^i\, \phi_j^i(T). \tag{6}$$

$c^i \in \mathbb{R}$ is the coefficient on the ith basis function. We solve for a new vector of c^i in each iteration j. To solve for N unknown scalars, we need at least N equations, which will be N copies of equation (6) evaluated at different "collocation" points on a grid that we construct in the state space. We can write these N equations in matrix form as

$$\mathbf{V_j} = \mathbf{\Phi_j}\, \mathbf{c_j},$$

where $\mathbf{V_j}$ and $\mathbf{c_j}$ are $N \times 1$ and $\phi_\mathbf{j}$ is $N \times N$. The vector of coefficients is solved by a matrix inversion:

$$\mathbf{c_j} = \mathbf{\Phi_j^{-1}}\, \mathbf{V_j}.$$

[31] In an abuse of notation, the subscript j now indexes the iteration, not time.

[32] The Stone-Weierstrass theorem tells us that every continuous function defined on a compact space can be approximated arbitrarily well by some polynomial.

In value function iteration, we assume some initial $\mathbf{c_0}$, maximize the right-hand side of the Bellman equation at the N collocation points, use the resulting vector of maximized values $\mathbf{V_1}$ to obtain a new vector of coefficients $\mathbf{c_1}$, and repeat until a predefined convergence criterion is satisfied. In order to verify the accuracy of our approximation, we should repeat the procedure with a different value of N or, nearly equivalently, vary the domain and check that the solution and policies are similar. As we vary the density of collocation points, the solution to the problem should not change by much. The appendix describes additional best practices.

Solving a finite-horizon problem works in a similar fashion. We begin with some terminal value function $\tilde{V}_\tau(T_\tau)$. If τ is sufficiently large and the policymaker discounts the future at some nonzero rate, then the choice of this value function will not matter for policy and welfare over some earlier, shorter horizon. The next step is to maximize the Bellman at time $\tau - 1$, which yields a vector of maximized values $\mathbf{V_{\tau-1}}$ that are used to recover the vector $\mathbf{c_{\tau-1}}$ and construct the time $\tau - 1$ value function approximant. We repeat the process, stepping back in time until time 0. At this point, we have a vector of coefficients (and thus a value function approximant) for every time period from 0 to τ and can therefore simulate policy trajectories from 0 to τ. As before, we should verify the accuracy of our solution by testing it with different values of N.

The final question to answer is: how do we select our basis functions and our set of points in the state space? The convention is to use Chebyshev polynomials for the basis functions and Chebyshev zeros for the points in the state space (after mapping the domain of the state space into the domain of the Chebyshev polynomials $[-1, 1]$) because of their favorable approximation properties when used together (Judd, 1998; Miranda and Fackler, 2002).[33] For each dimension d, we select a degree n_d Chebyshev polynomial and locate the n_d grid points at the n_d zeros of the polynomial. Boyd (2000, Chapter 4) describes the Chebyshev nodes as arising from spacing points evenly along a semicircle of unit radius and then mapping the points back onto the horizontal axis. This demonstration shows how Chebyshev nodes are unevenly spaced in the domain, tending to cluster near its edges. This spacing can be shown to work better than uniform spacing. To construct the full polynomial and grid, we take a tensor product of the unidimensional polynomials and grid points. Other approximating functions used in the literature include neural networks (Kelly and Kolstad, 1999b, 2001), which can be universal approximators of continuous functions on compact sets, and splines (Fitzpatrick and Kelly, 2017), which have advantages if the value function is nondifferentiable and which can target regions of the state space over which the value

[33] Boyd's Moral Principle (Boyd, 2000) essentially says to always use a Chebyshev basis, unless you are exceptionally sure that another basis is better for your problem.

function is especially curved.[34,35]

6.1 Frontier techniques

There are several drawbacks to the basic value function iteration approach outlined above. A first is that it is typically slow, particularly if the discount factor β is near 1, as is common in models with an annual timestep. A second is that the computational complexity of the problem increases exponentially in the dimension of the state space (commonly referred to as the Curse of Dimensionality). In the interest of advancing the literature, we here describe some ways to reduce the computational cost and/or improve the accuracy of the standard tensor product approximation approach described above.[36]

Recall that we created our grid of points in the state space with a tensor product. Our grid has n^d elements if we have n unique points for each state and d states. Adding more states to the model rapidly increases the number of grid points and thus the number of required Bellman maximizations. This computational cost has limited IAMs. Researchers have circumvented this problem by finding clever ways to reduce the dimensionality of a complex climate-economy model. For example, Kelly and Kolstad (1999b) and Traeger

[34] Traeger (2014) compares cubic spline and Chebyshev bases. Cubic splines are the most common choice of spline basis. These ensure continuous first and second derivatives, and they restrict the third derivative to be piecewise constant. However, we have seen that the third derivative of the value function plays a critical role in determining policy under uncertainty. The choice of cubic splines may therefore not be innocuous.

[35] Some recent papers have adopted nonstandard approaches to solving the Bellman equation, such as using "logarithmic", state-separable basis functions (Hwang et al., 2013; Hwang, 2016; Hwang et al., 2017) and fixing policy for some number of periods before approximating the remaining continuation value as a linear function of the per-period payoff (Heutel et al., 2015, 2016). We here caution against using less theoretically grounded methods for three reasons. First, several orders of value function derivatives and cross-derivatives matter for policy (see Section 3), and these alternative methods severely constrain these derivatives. A modeler should theoretically demonstrate that the correct value function indeed inherits the derivatives of per-period utility. Second, it is possible to "solve" a Bellman equation within a particular algorithm but for the solution to be incorrect. It is difficult to verify whether a solution to the approximated Bellman equation is correct, and theorems about function approximations provide extra confidence in the case of standard methods. Third, while some of these methods may work in particular cases of a particular model, they almost surely will not work in general cases. Yet the solutions can only be tested against known ones in particular (often deterministic) benchmark cases, even as a given paper is interested in the difference between a benchmark case and some other case. It is hard to know whether the cases of interest in a given paper are also particular cases in which the alternative algorithm happens to work and to know whether the reported differences in, for example, policy between these cases and the benchmark case are merely due to the numerical methods.

[36] Many techniques common in other settings are often not helpful in solving integrated assessment models. Standard Euler equation methods often do not help because IAMs contain more states than controls. One may also attempt to use shape-preserving methods (Judd and Solnick, 1994; Cai and Judd, 2012), but they are often not justified because we do not know in advance that the value function is, for example, concave, and we may even know that it is not concave everywhere (due to nonlinearities in the forcing equation, for example).

(2014) employ models with one- and two-state climate systems instead of the five-state climate system of the DICE model (which is already a coarse approximation to full climate models).[37]

There are ways to reduce computation time without sacrificing model complexity. A simple way is to supply the solver used in the maximization step with gradient and Hessian information.[38] Most solvers allow the user to supply analytic gradients and Hessians so that the solver no longer needs to approximate them, thereby improving both accuracy and speed. However, solving for these functions by hand, particularly in complex models like IAMs, can be troublesome and prone to human error. Autodifferentiation is a powerful technique that exploits the arithmetic structure of computer code to automatically generate gradients and Hessians. Since all programs are a combination of basic operations such as addition or exponentiation, we can use a prepackaged autodifferentiation algorithm to repeatedly apply the chain rule to arbitrary computer code in order to take derivatives of any program.

For example, consider maximizing the right-hand side of a Bellman equation. Many models (including DICE) use an isoelastic per-period utility function. In this case, the utility function has a closed-form derivative that can be easily computed by an algorithm ($\frac{d}{dc_t} c_t^{1-\eta}/(1-\eta) = c_t^{-\eta}$). The value function approximant is often just a polynomial of the next period's states, and the next period's states are generally simple functions of the current period's actions and states. So the derivative of the continuation value with respect to the current control can be computed by applying the chain rule to the transition equations and the value function approximant. The modeler does not need to solve for the analytic gradient at any point because the autodifferentiation algorithm can recognize the functional forms in the code that make up the right-hand side of the Bellman and thus can apply derivatives and the chain rule to back out the gradient. Autodifferentiation can greatly reduce the potential for error in hand-coding these functions. Note, however, that the optimization step is not computationally challenging in many recursive IAMs, so the gains from supplying gradients and Hessians may be small in many applications.

For infinite-horizon problems, we can accelerate convergence of the value function iteration algorithm outlined above by using a technique called "modified policy iteration". Modified policy iteration exploits two facts. First, maximizing the Bellman equation is the costly step in the algorithm. Second, the optimal policy function corresponding to the Bellman in iteration j will be very close to the optimal policy function for the Bellman in iteration $j+1$, especially on the grid of interest. Instead of maximizing the Bellman in each

[37] These papers solve an infinite-horizon model with fewer states than the full model. Traeger (2014) shows that his transition equations may fit the desired climate dynamics better than does the larger set of transition equations in the full DICE model. Lemoine and Traeger (2014) approximate the omitted states as functions of the tracked states. Other papers have used finite-horizon methods that do not require iterating over the full grid of nodes as many times (e.g., Cai et al., 2012). Recent work has used sparse grids (discussed below) to solve a finite-horizon version of a full IAM (Rudik, 2016).

[38] Economists could also take greater advantage of high-performance computing facilities (Dongarra and van der Steen, 2012).

iteration to recover a new optimal policy, modified policy iteration periodically reuses the previous iteration's optimal policy to get the "maximized" Bellman value and solve for a new vector of coefficients for iteration $j+2$.[39] Puterman and Shin (1978) demonstrate that this method will converge at least as fast as standard value function iteration.

The domain of the value function approximant strongly affects both speed and accuracy. Selecting a domain that is larger than necessary will result in worse accuracy for a given grid, but selecting a domain that is too small can lead the solver to evaluate states that lie outside the domain of the approximant, where accuracy deteriorates. A general rule of thumb is to make the domain only as large as it needs to be to ensure that relevant areas of the state space do not transition outside the domain and to ensure that calculating expectations does not require placing too much weight on points outside the domain. If the domain is any larger, then we are fitting the value function in regions of the state space where we may never travel in simulations, which increases the computational burden with no benefit.

Some techniques adapt the domain of the value function approximant to the domain of simulated trajectories. The domain of the standard tensor product grid is shaped as a hypercube, but the domain of the simulated trajectories may not be. For example, they may be clustered along a diagonal of the grid in a three-dimensional problem. In this case, large areas of the domain may never be reached in the simulations. Placing grid points in these areas adds computational cost with little benefit in improving the accuracy of our simulated outcomes. Adaptive grids may be able to reduce computational cost without sacrificing accuracy.

Once we obtain the set of simulated trajectories from an existing solution to our model, perhaps from a version that uses a hypercube-shaped grid, we can adapt the domain of the value function approximant by using a change of coordinates. Judd et al. (2014b) and Maliar and Maliar (2015) propose a principal component transformation of the simulated data to rotate the coordinate system of the domain. When a hypercube is fit around the simulated points in the principal components system, the grid points are packed tightly into the important regions of the state space, allowing the modeler to obtain an approximation of the same accuracy but with fewer grid points.

In finite-horizon problems, a modeler can construct a time-adaptive grid. In the example of our simple IAM presented earlier, the only area of the state space we are concerned about at the initial time $t=0$ is T_0. The domain of the value function at time $t=0$ can be tightly bound around that value. Moving forward to time $t=1$, the transition equations govern the maximum and minimum possible temperature, which tells us how big the domain could possibly need to be conditional on the variance of the temperature shocks. In general, the domain becomes larger as we move forward in time because sequences of shocks may

[39] Generally, the closer the value function approximant is to the true value function, the less change there will be in the optimal policy conditional on the value function approximant from iteration to iteration. Therefore, as the value function iteration algorithm progresses, the "optimal" policy can be used an increasing number of times between each value function maximization step.

reinforce each other.

Economists have begun to explore sparse grids (Smolyak, 1963; Krueger and Kubler, 2004; Judd et al., 2014b; Rudik, 2016). Sparse grids are constructed similarly to the standard tensor product grid; however, they omit certain subspaces of the full tensor product grid that are less important for approximation quality. The benefit of sparse grids is that the number of grid points only increases polynomially in the dimensionality of the problem. Although they use fewer grid points, sparse grids can be very accurate (Barthelmann et al., 2000). Indeed, one can think of sparse grids as the solution to an optimization problem: for an exogenously given number of nodes, find the approximation space that yields the greatest accuracy in terms of the L_2 and L_∞ norms within the space of functions with bounded second-order mixed derivatives (Brumm and Scheidegger, 2016).[40] Sparse grids promise significant computational savings relative to standard methods.

7 The Second Wave of Research with Recursive Models

The first wave of research with recursive IAMs devoted substantial efforts to solving models with a single type of uncertainty. We have suggested that frontier techniques can open up a broader range of modeling possibilities. Recent work has begun applying such techniques to ask questions about model uncertainty (Rudik, 2016). We here suggest some potentially helpful directions for a second wave of work with recursive IAMs, most of which depend on the ability to solve models with dramatically larger state spaces than used in the first wave of research.

1. Recursive IAMs should better model learning. Standard IAMs dramatically simplify the evolution of the climate and the economy, and recursive IAMs dramatically simplify the evolution of knowledge. This simplification is potentially critical because, as implied by Section 3, assuming that learning occurs too rapidly or that additional emissions generate too much information can qualitatively change policy conclusions. Actual knowledge about, for instance, future warming has been limited by the incompleteness of scientific observations, by the combination of complexity and coarseness in numerical climate models, and by the presence of multiple, potentially correlated uncertain parameters (e.g., climate sensitivity, ocean heat content, and aerosol forcing). Future work with recursive IAMs should seek not just to model the climate system but to model those features of the climate system that make climate science hard. Such work can answer important, but often overlooked, policy questions about the value of scientific monitoring and the optimal types of monitoring. As but one example,

[40]Most finite-horizon versions of the DICE IAM satisfy the bounded mixed derivative condition.

allowing for hidden states could be critical to asking questions about the value of an ocean monitoring system.

2. Economic uncertainties are more severe than scientific uncertainties, and are potentially as important for policy. Looking out 100 years, the overall quality of technology, the level of consumption, the quality of abatement technology, and preferences over environmental quality are at least as uncertain as any scientific aspect of the climate. Such uncertainties are very difficult to adequately calibrate or model, but doing so seems critical. Some of these uncertainties could be calibrated from asset prices, and others could be calibrated from historical time series.

3. Some of the most pressing unknowns relate to the possibility of catastrophes. In particular, many scientists and economists worry about the potential for abrupt, irreversible changes in the climate system ("tipping points") and about the potential for large declines in environmental quality or consumption in response to future warming. Thus far, work on physical tipping points has used reduced-form representations of the tipping process, and work on catastrophes has used either permanent shocks to output or reduced-form uncertainty about the curvature of the damage function. Future work should attempt a more structured approach to each problem, which may require extending IAMs to more faithfully represent the actual sources of information about these possible events as well as the actual, nonlinear mechanisms through which these events would unfold.

4. Future work should consider moving beyond the DICE model. To date, perhaps all recursive IAMs either have used variants of the DICE model (Nordhaus, 1992, 2008) or have used still simpler models. The DICE model has been an enormously valuable benchmark, and its transparent structure has been a boon to qualitative understanding. We showed above how to decompose the value function to glean theoretical insight from recursive settings. This type of approach to recursive modeling promises qualitative insight even from settings more complicated than DICE. Further, policymakers are in fact using IAMs for their quantitative policy conclusions (Greenstone et al., 2013). It is important to know how greater realism affects these numbers. Future advances in IAMs should consider extending benchmark settings to allow for features like multisector economies, endogenous or directed forms of growth, more realistic scientific modules, and a treatment of space. But this effort should be disciplined. The community should first aim to develop a set of stylized facts about asset prices, climate science, and economic activity that integrated assessment models should replicate.

The title of this review refers to recursive integrated assessment as being at an inflection point. Seven years ago, the value of recursive modeling was convex in modeling effort, but this value appears to have become concave more recently. We hope that the research directions suggested here will make that value convex once again. We see the emerging potential to integrate next-level considerations into full IAMs and to ask new questions that take the information structure of climate change as seriously as the best-guess physical dynamics. We encourage recursive modelers to exercise creativity in moving towards these and other new directions.

References

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.

Anthoff, David and Richard S. J. Tol (2013) "The uncertainty about the social cost of carbon: A decomposition analysis using FUND," *Climatic Change*, Vol. 117, No. 3, pp. 515–530.

Anthoff, David, Richard S. J. Tol, and Gary W. Yohe (2009) "Risk aversion, time preference, and the social cost of carbon," *Environmental Research Letters*, Vol. 4, No. 2, p. 024002.

Baker, Marcia B. and Gerard H. Roe (2009) "The shape of things to come: Why is climate change so predictable?" *Journal of Climate*, Vol. 22, No. 17, pp. 4574–4589.

Barthelmann, Volker, Erich Novak, and Klaus Ritter (2000) "High dimensional polynomial interpolation on sparse grids," *Advances in Computational Mathematics*, Vol. 12, No. 4, pp. 273–288.

Becker, Gary S., Kevin M. Murphy, and Robert H. Topel (2010) "On the economics of climate policy," *The B.E. Journal of Economic Analysis & Policy*, Vol. 10, No. 2, p. Article 19.

Boyd, John P. (2000) *Chebyshev and Fourier Spectral Methods*, Mineola, New York: Dover Publications, 2nd edition.

Breeden, Douglas T. (1979) "An intertemporal asset pricing model with stochastic consumption and investment opportunities," *Journal of Financial Economics*, Vol. 7, No. 3, pp. 265–296.

Brumm, Johannes and Simon Scheidegger (2016) "Using adaptive sparse grids to solve high-dimensional dynamic models," working Paper.

Cai, Yongyang and Kenneth L. Judd (2012) "Dynamic programming with shape-preserving rational spline Hermite interpolation," *Economics Letters*, Vol. 117, No. 1, pp. 161–164.

——— (2014) "Advances in numerical dynamic programming and new applications," in Karl Schmedders and Kenneth L. Judd eds. *Handbook of Computational Economics*, Vol. 3: Elsevier, pp. 479–516, doi: 10.1016/B978-0-444-52980-0.00008-6.

Cai, Yongyang, Kenneth L. Judd, Timothy M. Lenton, Thomas S. Lontzek, and Daiju Narita (2015) "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 (2012) "DSICE: A dynamic stochastic integrated model of climate and economy," Working Paper 12-02, The Center for Robust Decision Making on Climate and Energy Policy.

——— (2013) "The social cost of stochastic and irreversible climate change," Working Paper 18704, National Bureau of Economic Research.

Cai, Yongyang, Timothy M. Lenton, and Thomas S. Lontzek (2016) "Risk of multiple interacting tipping points should encourage rapid CO2 emission reduction," *Nature Climate Change*, Vol. 6, No. 5, pp. 520–525.

Carroll, Christopher D. and Miles S. Kimball (1996) "On the concavity of the consumption function," *Econometrica*, Vol. 64, No. 4, pp. 981–992.

Crainich, David, Louis Eeckhoudt, and Alain Trannoy (2013) "Even (mixed) risk lovers are prudent," *American Economic Review*, Vol. 103, No. 4, pp. 1529–1535.

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.

Cyert, Richard M. and Morris H. DeGroot (1974) "Rational expectations and bayesian analysis," *Journal of Political Economy*, Vol. 82, No. 3, pp. 521–536.

Dietz, Simon, Christian Gollier, and Louise Kessler (2015) "The climate beta," Working Paper 215, Centre for Climate Change Economics and Policy, Grantham Institute on Climate Change and the Environment.

Dietz, Simon and Nicholas Stern (2015) "Endogenous growth, convexity of damages and climate risk: How Nordhaus' framework supports deep cuts in carbon emissions," *The Economic Journal*, Vol. 125, No. 583, pp. 574–620, Forthcoming.

Dongarra, J. J. and A. J. van der Steen (2012) "High-performance computing systems: Status and outlook," *Acta Numerica*, Vol. 21, pp. 379–474.

Drèze, Jacques H. and Franco Modigliani (1972) "Consumption decisions under uncertainty," *Journal of Economic Theory*, Vol. 5, No. 3, pp. 308–335.

Eeckhoudt, Louis, Christian Gollier, and Thierry Schneider (1995) "Risk-aversion, prudence and temperance: A unified approach," *Economics Letters*, Vol. 48, No. 34, pp. 331–336.

Epstein, Larry G. and Stanley E. Zin (1989) "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, Vol. 57, No. 4, pp. 937–969.

Fernández-Villaverde, Jesús, Juan F. Rubio Ramírez, and Frank Schorfheide (2016) "Solution and estimation methods for DSGE models," Working Paper 21862, National Bureau of Economic Research.

Fisher, Anthony C. and Urvashi Narain (2003) "Global warming, endogenous risk, and irreversibility," *Environmental and Resource Economics*, Vol. 25, No. 4, pp. 395–416.

Fitzpatrick, Luke G. and David L. Kelly (2017) "Probabilistic stabilization targets," *Journal of the Association of Environmental and Resource Economists*, Vol. forthcoming.

Gillingham, Kenneth, William D. Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, John Reilly, and Paul Sztorc (2015) "Modeling uncertainty in climate change: A multi-model comparison," Working Paper 21637, National Bureau of Economic Research.

Gollier, Christian (2001) *The Economics of Risk and Time*, Cambridge, Mass: MIT Press.

——— (2010) "Ecological discounting," *Journal of Economic Theory*, Vol. 145, No. 2, pp. 812–829.

——— (2012) "Evaluation of long-dated investments under uncertain growth trend, volatility and catastrophes," CESifo Working Paper 4052.

Gollier, Christian, Bruno Jullien, and Nicolas Treich (2000) "Scientific progress and irreversibility: an economic interpretation of the 'Precautionary Principle'," *Journal of Public Economics*, Vol. 75, No. 2, pp. 229–253.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014) "Optimal taxes on fossil fuel in general equilibrium," *Econometrica*, Vol. 82, No. 1, pp. 41–88.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton (2011) "Estimating the social cost of carbon for use in U.S. federal rulemakings: a summary and interpretation," *National Bureau of Economic Research Working Paper Series*, Vol. No. 16913.

——— (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.

Hansen, J., G. Russell, A. Lacis, I. Fung, D. Rind, and P. Stone (1985) "Climate response times: Dependence on climate sensitivity and ocean mixing," *Science*, Vol. 229, No. 4716, pp. 857–859.

Heutel, Garth, Juan Moreno Cruz, and Soheil Shayegh (2015) "Solar geoengineering, uncertainty, and the price of carbon," Working Paper 21355, National Bureau of Economic Research.

——— (2016) "Climate tipping points and solar geoengineering," *Journal of Economic Behavior and Organization*, Vol. 132, pp. 19–45.

Hope, Chris (2006) "The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern," *The Integrated Assessment Journal*, Vol. 6, No. 1, pp. 19–56.

Howarth, Richard B. (2003) "Discounting and uncertainty in climate change policy analysis," *Land Economics*, Vol. 79, No. 3, pp. 369–381.

Hwang, In Chang (2016) "A recursive method for solving a climate-economy model: Value function iterations with logarithmic approximations," *Computational Economics*, pp. 1–16.

Hwang, In Chang, Frédéric Reynès, and Richard S. J. Tol (2017) "The effect of learning on climate policy under fat-tailed uncertainty," *Resource and Energy Economics*, Vol. 48, pp. 1–18.

Hwang, In Chang, Richard Tol, and Marjan Hofkes (2013) "Active learning about climate change," Working Paper 65-2013, Department of Economics, University of Sussex.

Jensen, Svenn and Christian Traeger (2016) "Pricing climate risk," working paper.

Jensen, Svenn and Christian P. Traeger (2014) "Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings," *European Economic Review*, Vol. 69, pp. 104–125.

Jones, Robert A. and Joseph M. Ostroy (1984) "Flexibility and uncertainty," *The Review of Economic Studies*, Vol. 51, No. 1, pp. 13–32.

Judd, Kenneth L. (1998) *Numerical Methods in Economics*, Cambridge, Mass.: MIT Press.

Judd, Kenneth L., Lilia Maliar, and Serguei Maliar (2014a) "Lower bounds on approximation errors: Testing the hypothesis that a numerical solution is accurate," working paper.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero (2014b) "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.

Judd, Kenneth L. and Andrew Solnick (1994) "Numerical dynamic programming with shape-preserving splines."

Kelly, David L. and Charles D. Kolstad (1999a) "Integrated assessment models for climate change control," in Henk Folmer and Tom Tietenberg eds. *International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues*, Cheltenham, UK: Edward Elgar, pp. 171–197.

——— (1999b) "Bayesian learning, growth, and pollution," *Journal of Economic Dynamics and Control*, Vol. 23, No. 4, pp. 491–518.

——— (2001) "Solving infinite horizon growth models with an environmental sector," *Computational Economics*, Vol. 18, No. 2, pp. 217–231.

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, No. 2012-13.

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 J. (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 (2016) "The climate risk premium: How uncertainty affects the social cost of carbon," Working Paper 15-01, University of Arizona.

Lemoine, Derek M. (2010) "Climate sensitivity distributions depend on the possibility that models share biases," *Journal of Climate*, Vol. 23, No. 16, pp. 4395–4415.

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.

Lemoine, Derek and Christian P. Traeger (2016a) "Ambiguous tipping points," *Journal of Economic Behavior & Organization*, Vol. 132, pp. 5–18.

——— (2016b) "Economics of tipping the climate dominoes," *Nature Climate Change*, Vol. 6, No. 5, pp. 514–519.

Litterman, Bob (2013) "What is the right price for carbon emissions?" *Regulation*, Vol. 36, No. 2, pp. 38–43.

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, pp. 441–444.

Lucas, Robert E., Jr. (1978) "Asset prices in an exchange economy," *Econometrica*, Vol. 46, No. 6, pp. 1429–1445.

Maliar, Lilia and Serguei Maliar (2014) "Numerical methods for large-scale dynamic economic models," in Karl Schmedders and Kenneth L. Judd eds. *Handbook of Computational Economics*, Vol. 3: Elsevier, pp. 325–477, doi: 10.1016/B978-0-444-52980-0.00007-4.

——— (2015) "Merging simulation and projection approaches to solve high-dimensional problems with an application to a new Keynesian model," *Quantitative Economics*, Vol. 6, No. 1, pp. 1–47.

Manne, Alan and Richard Richels (1995) "The greenhouse debate: Economic efficiency, burden sharing and hedging strategies," *The Energy Journal*, Vol. 16, No. 4, pp. 1–37.

Miranda, Mario J. and Paul L. Fackler (2002) *Applied Computational Economics and Finance*, Cambridge, Massachusetts: MIT Press.

Nordhaus, William (2013) "Integrated economic and climate modeling," in Peter B. Dixon and Dale W. Jorgenson eds. *Handbook of Computable General Equilibrium Modeling*, Vol. 1: North Holland, Elsevier B.V. pp. 1069–1131.

Nordhaus, William D. (1992) "An optimal transition path for controlling greenhouse gases," *Science*, Vol. 258, No. 5086, pp. 1315–1319.

——— (1994) *Managing the Global Commons: The Economics of Climate Change*, Cambridge, Mass: The MIT Press, 1st edition.

——— (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven: Yale University Press.

——— (2016) "Projections and uncertainties about climate change in an era of minimal climate policies," Working Paper 22933, National Bureau of Economic Research.

Nordhaus, William D. and David Popp (1997) "What is the value of scientific knowledge? An application to global warming using the PRICE model," *Energy Journal*, Vol. 18, No. 1, pp. 1–45.

Peck, Stephen C. and Thomas J. Teisberg (1993) "Global warming uncertainties and the value of information: an analysis using CETA," *Resource and Energy Economics*, Vol. 15, No. 1, pp. 71–97.

Pizer, William A. (1999) "The optimal choice of climate change policy in the presence of uncertainty," *Resource and Energy Economics*, Vol. 21, No. 34, pp. 255–287.

Puterman, Martin L. and Moon Chirl Shin (1978) "Modified policy iteration algorithms for discounted Markov decision problems," *Management Science*, Vol. 24, No. 11, pp. 1127–1137.

Pycroft, Jonathan, Lucia Vergano, Chris W. Hope, Daniele Paci, and Juan Carlos Ciscar (2011) "A tale of tails: Uncertainty and the social cost of carbon dioxide," *Economics: The Open-Access, Open-Assessment E-Journal*, Vol. 5, No. 2011-22.

Raper, Sarah C. B., Jonathan M. Gregory, and Ronald J. Stouffer (2002) "The role of climate sensitivity and ocean heat uptake on AOGCM transient temperature response," *Journal of Climate*, Vol. 15, No. 1, pp. 124–130.

Roe, Gerard H. (2009) "Feedbacks, timescales, and seeing red," *Annual Review of Earth and Planetary Sciences*, Vol. 37, No. 1, pp. 93–115.

Roe, Gerard H. and Marcia B. Baker (2007) "Why is climate sensitivity so unpredictable?" *Science*, Vol. 318, No. 5850, pp. 629–632.

Rudik, Ivan (2016) "Optimal climate policy when damages are unknown," working paper, Iowa State University.

Sandsmark, Maria and Haakon Vennemo (2007) "A portfolio approach to climate investments: CAPM and endogenous risk," *Environmental and Resource Economics*, Vol. 37, No. 4, pp. 681–695.

Santos, Manuel S. (2000) "Accuracy of numerical solutions using the Euler equation residuals," *Econometrica*, Vol. 68, No. 6, pp. 1377–1402.

Sibley, David S. (1975) "Permanent and transitory income effects in a model of optimal consumption with wage income uncertainty," *Journal of Economic Theory*, Vol. 11, No. 1, pp. 68–82.

Smolyak, S. (1963) "Quadrature and interpolation formulas for tensor products of certain classes of functions," *Soviet Mathematics, Doklady*, Vol. 4, pp. 240–243.

Stokey, Nancy and Robert E. Lucas, Jr. (1989) *Recursive Methods in Economic Dynamics*, Cambridge Mass.: Harvard University Press.

Tol, Richard S. J. (1999) "Safe policies in an uncertain climate: an application of FUND," *Global Environmental Change*, Vol. 9, No. 3, pp. 221–232.

Traeger, Christian (2014) "A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change," *Environmental and Resource Economics*, Vol. 59, No. 1, pp. 1–37.

Traeger, Christian P. (2015) "Analytic integrated assessment and uncertainty," working paper.

Urban, Nathan M. and Klaus Keller (2009) "Complementary observational constraints on climate sensitivity," *Geophysical Research Letters*, Vol. 36, p. L04708.

Weil, Philippe (1989) "The equity premium puzzle and the risk-free rate puzzle," *Journal of Monetary Economics*, Vol. 24, No. 3, pp. 401–421.

Weitzman, Martin L. (2013) "Tail-hedge discounting and the social cost of carbon," *Journal of Economic Literature*, Vol. 51, No. 3, pp. 873–882.

Appendices

The first appendix extends our theoretical analysis to the case of Epstein-Zin-Weil preferences. The second appendix describes our implementation of the DICE integrated assessment model and demonstrates some best practices.

Appendix A: Extension to Epstein-Zin-Weil Preferences

Epstein-Zin-Weil preferences (also called "recursive utility") allow for separation between preferences for smoothing consumption over risk and over time: they disentangle risk aversion from the elasticity of intertemporal substitution. Let ψ be the Arrow-Pratt measure of relative risk aversion and η be the reciprocal of the intertemporal elasticity of substitution. With Epstein-Zin-Weil preferences, the main text's Bellman equation becomes

$$V_t(T_t, \mu_t, \Sigma_t) = \max_{e_t} \left\{ \frac{c_t(e_t;T_t)^{1-\eta}}{1-\eta} + \frac{\beta}{1-\eta} \left(E\left[\left((1-\eta)\, V_{t+1}(T_{t+1}, \mu_{t+1}, \Sigma_{t+1}) \right)^{\frac{1-\psi}{1-\eta}} \right] \right)^{\frac{1-\eta}{1-\psi}} \right\},$$

for some consumption function $c_t(e_t;T_t)$. Expected utility corresponds to the special case of $\psi = \eta$. The first-order condition becomes

$$\begin{aligned} c_t^{-\eta} \frac{\partial c_t}{\partial e_t} = &- \beta \left(E_t\left[((1-\eta)V_{t+1})^{\frac{1-\psi}{1-\eta}} \right] \right)^{\frac{\psi-\eta}{1-\psi}} \\ & E_t\left[((1-\eta)V_{t+1})^{\frac{\eta-\psi}{1-\eta}} \left(\frac{\partial V_{t+1}}{\partial T_{t+1}} \frac{\partial T_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \mu_{t+1}} \frac{\partial \mu_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \Sigma_{t+1}} \frac{\partial \Sigma_{t+1}}{\partial e_t} \right) \right]. \end{aligned}$$

After a second-order Taylor expansion of the second term, the right-hand side becomes:

$$\begin{aligned} &- \beta \left(E_t\left[((1-\eta)V_{t+1})^{\frac{1-\psi}{1-\eta}} \right] \right)^{\frac{\psi-\eta}{1-\psi}} E_t\left[((1-\eta)V_{t+1})^{\frac{\eta-\psi}{1-\eta}} \right] \overbrace{E_t\left[\frac{\partial V_{t+1}}{\partial T_{t+1}} \frac{\partial T_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \mu_{t+1}} \frac{\partial \mu_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \Sigma_{t+1}} \frac{\partial \Sigma_{t+1}}{\partial e_t} \right]}^{\text{channels from main text}} \\ &+ \underbrace{\beta \left(E_t\left[((1-\eta)V_{t+1})^{\frac{1-\psi}{1-\eta}} \right] \right)^{\frac{\psi-\eta}{1-\psi}} Cov_t\left[((1-\eta)V_{t+1})^{\frac{\eta-\psi}{1-\eta}},\ - \left(\frac{\partial V_{t+1}}{\partial T_{t+1}} \frac{\partial T_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \mu_{t+1}} \frac{\partial \mu_{t+1}}{\partial e_t} + \frac{\partial V_{t+1}}{\partial \Sigma_{t+1}} \frac{\partial \Sigma_{t+1}}{\partial e_t} \right) \right]}_{\text{preference for temporal resolution}} \end{aligned}$$

The first line is as analyzed in the main text, with all channels adjusted by the factor

$$\left(E_t\left[((1-\eta)V_{t+1})^{\frac{1-\psi}{1-\eta}} \right] \right)^{\frac{\psi-\eta}{1-\psi}} E_t\left[((1-\eta)V_{t+1})^{\frac{\eta-\psi}{1-\eta}} \right].$$

Note that this factor is 1 in the case with $\psi = \eta$, which corresponds to the case of expected utility analyzed in the main text. The second line is novel. It is zero when $\psi = \eta$, which explains why it is absent from the main text's analysis of a setting with expected utility. When $\psi > \eta$, the policymaker displays a preference for early resolution of uncertainty: she would pay money just to obtain information about a future risk sooner, even when she cannot act on that information. The covariance captures how the marginal value of emission reductions covaries with welfare. The second line increases the optimal emission tax if and only if this covariance is positive. When $\psi > \eta$ (as in standard calibrations), the covariance is positive when large V_{t+1} pairs with a small marginal value of emission reductions. For example, the covariance is positive when small values for climate sensitivity imply both high welfare and a small marginal value of emission reductions. We might in fact expect this combination to be the case, in which case the preference for temporal resolution channel increases the optimal emission tax. Intuitively, when the covariance is positive, reducing emissions smooths future welfare outcomes because it increases welfare most strongly in states with low welfare. Additional emission reductions thus help to resolve uncertainty about future welfare at an earlier date, and a policymaker with $\psi > \eta$ will pay for this earlier resolution of uncertainty.

Figure A-1 plots the quantitatively important channels in this setting with Epstein-Zin-Weil preferences. The left panels keep relative risk aversion at its DICE value of 2 but increase the elasticity of intertemporal substitution to 2/3, and the right panels keep the elasticity of intertemporal substitution at its DICE value of 1/2 but increase relative risk aversion to 3. These changes are in the same direction as changes suggested by some recent asset pricing models. The top panels plot the case of persistently uncertain climate sensitivity, and the bottom panels allow the policymaker to learn about climate sensitivity from observations of temperature (and to anticipate that she will do so). The active learning channel is constructed in the same way as in the main text's case of expected utility. Raising relative risk aversion does not affect the certainty-equivalent emission tax, but by reducing the consumption discount rate, raising the elasticity of intertemporal substitution increases the certainty-equivalent tax by over 40% (around \$3.25 per tCO_2) in the first period. In all cases, the temporal resolution channel is positive (as expected) but very small, in part because our experiments do not change preferences by a large amount. The other channels are qualitatively similar to those in the main text's case of expected utility. We note three differences. First, by making the policymaker more sensitive to the risk induced by future revisions to beliefs, raising relative risk aversion increases the size of the signal smoothing channel in early years and leads it to decline monotonically over time (at least in this simulation, where random variables happen to take on their mean values). Second, raising relative risk aversion reduces the uncertainty adjustment by a bit. Third, raising relative risk aversion also makes the mean-belief component of the active learning channel start off negative, which makes active learning reduce the first period's optimal emission tax. These last two effects are a bit of a puzzle, but they could be related to how a preference for early resolution of uncertainty makes the policymaker dislike persistent shocks to mean beliefs, which are especially likely early on and, owing to interactions with inertia, especially likely when climate sensitivity is small.

Appendix B: The DICE Model

We now describe the model equations and calibration. We then extend our theoretical analysis of the channels through which uncertainty affects policy to the full numerical setting. We conclude with best practices for validating recursive models.

The DICE model is a Ramsey growth model coupled to a climate module. An infinitely lived representative agent aims to maximize the sum of the stream of discounted utility from consuming output. In decade t, the agent begins with some level of capital K_t. To produce gross output Y_t^g, the agent combines capital with labor L_t and technology A_t in a Cobb-Douglas production function:

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

Some of this output is lost to damages caused by surface warming T_t^s, so that output net of damages is given by

$$Y_t^n = \frac{Y_t^g}{1 + b_2\ [T_t^s]^{b_3}}.$$

The agent can allocate her net output to consumption C_t or emissions abatement α_t, with residual output invested towards future capital, which depreciates at an annual rate δ_k. Per-period utility is:

$$u(C_t; L_t) = L_t \frac{(C_t/L_t)^{1-\eta}}{1-\eta},$$

with $\eta \geq 0, \neq 1$. Emissions e_t (net of abatement) are

$$e_t = 10 * \left[\sigma_t (1-\alpha_t) Y_t^g + B_t\right],$$

where B_t gives exogenous emissions at time t and σ_t is the emission intensity of production at time t. We will constrain abatement to be less than the emissions generated by factor production. Emissions enter the atmospheric CO_2 stock M_t^{atm}. Between decades, CO_2 mixes between the atmospheric reservoir and reservoirs in the upper (M_t^{up}) and lower (M_t^{lo}) ocean.

Additional atmospheric CO_2 increases radiative forcing $F_t(M_t^{atm})$, which measures energy at the earth's surface. Forcing is given by

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

where EF_t is exogenous forcing from other long-lived greenhouse gases and f_{2x} is the amount of forcing from doubling CO_2. Heat transfers between two reservoirs: one reservoir T_t^s that

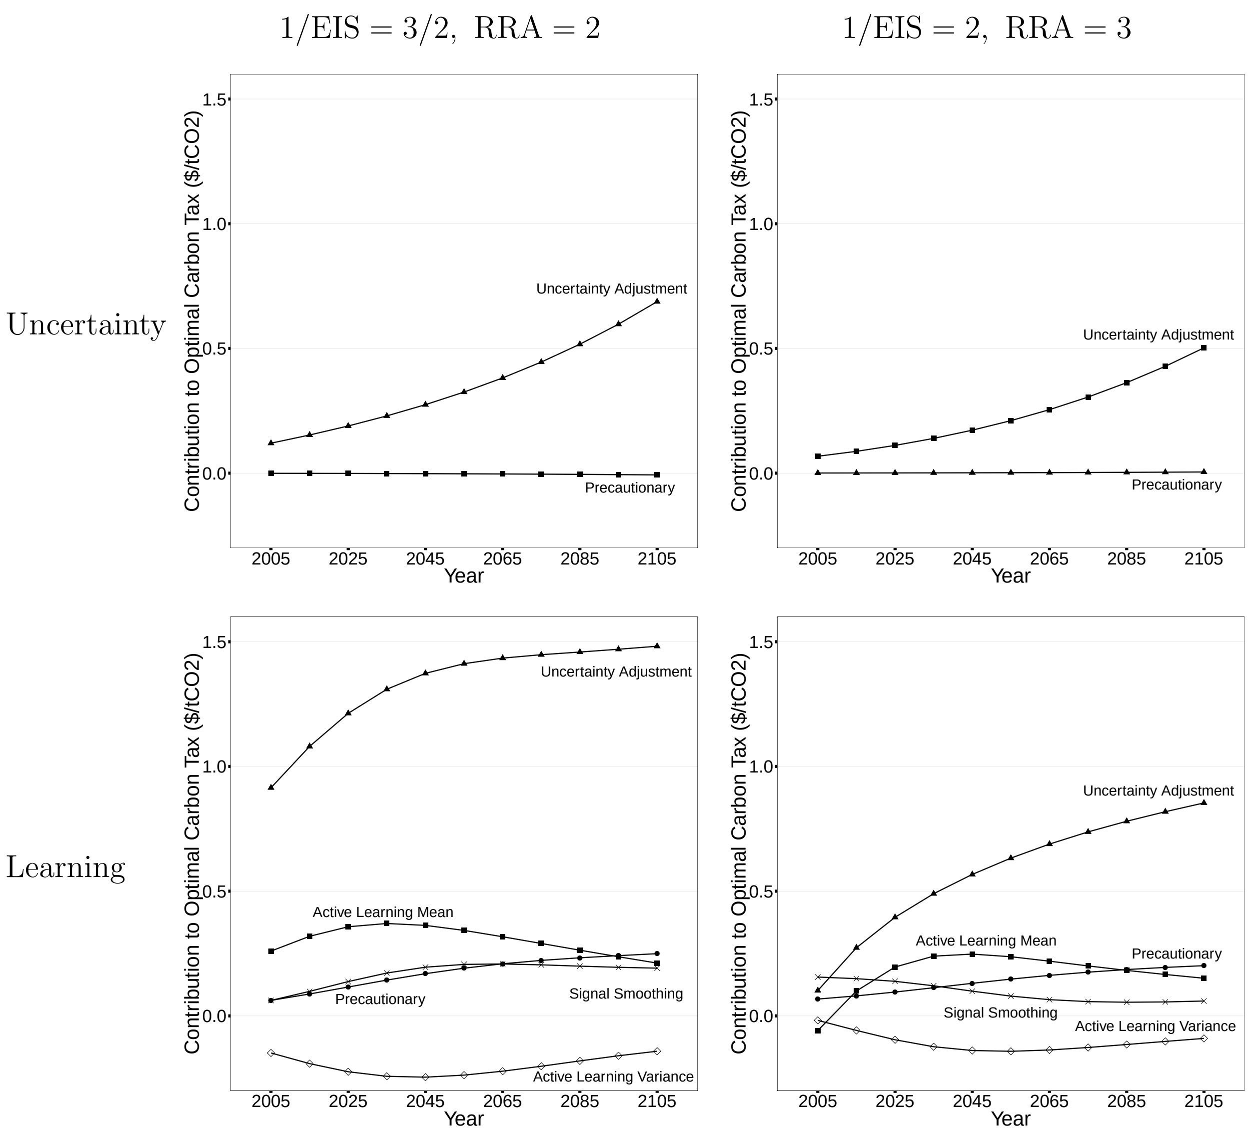

Figure A-1: The channels that determine how uncertainty about the climates sensitivity to emissions changes the optimal emission tax, for settings without learning (top row), with anticipated learning (bottom row), with less desire to smooth consumption over time (left column), and with greater aversion to risk (right column). All simulations fix random variables at their mean values.

includes the earth's surface and the upper ocean, and a second reservoir T_t^o that reflects the lower ocean. Each reservoir's rate of warming is determined by three parameters: C_3 governs the heat loss from the surface to the lower ocean, C_4 governs the heat loss from the lower ocean to the surface, and C_1 governs how quickly surface temperature responds to changes in forcing or to differences in the surface-ocean temperature gradient.

The climate sensitivity s determines how much the earth eventually warms after doubling of CO_2. We model the climate sensitivity as in Roe and Baker (2007):

$$s = \frac{\lambda_0}{1-\Delta},$$

where λ_0 is the climate sensitivity for a reference system (lacking earth system feedbacks) and $\Delta < 1$ is the climate system's feedback factor. See Roe (2009) and Lemoine (2010) for further exposition of linear feedback analysis.

To admit a closed-form updating rule for beliefs, we take Δ as the unknown parameter instead of s and, following Roe and Baker (2007), we assume that the agent believes $\Delta \sim \mathcal{N}(\mu_t, \Sigma_t)$ at time t. Her ability to learn about the true value of Δ is hindered by an independent and identically distributed temperature shock $\epsilon_t \sim \mathcal{N}(0, \sigma_T^2)$. Each decade, the agent updates the mean μ_t and variance Σ_t of her beliefs over Δ by using a normal-normal conjugate updating rule. The variance of her beliefs declines deterministically and monotonically over time.[41] However the mean of her belief at time $t+1$ is a function of the unknown temperature in time $t+1$, so μ_{t+1} is unknown at time t and evolves stochastically.

The model's exogenous economic processes are

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

where index $t = 0$ indicates the year 2005, $t = 1$ indicates the year 2015, and so on. The model's exogenous climate-related processes are

$$\begin{aligned}
\sigma_t &= \sigma_{t-1}/(1 - g_{\sigma,t}) && \text{(Gross emissions per unit of output)}\\
g_{\sigma,t} &= g_{\sigma,0} \exp(-\delta_\sigma\, t) && \text{(Growth rate of gross emissions per unit of output)}\\
\psi_t &= \frac{a_0\, \sigma_t}{a_1 a_2}\left(a_1 - 1 + \exp(-g_\Psi\, t)\right) && \text{(Abatement cost coefficient)}\\
B_t &= B_0\, g_B^t && \text{(Non-industrial } CO_2 \text{ emissions)}\\
EF_t &= EF_0 + 0.1\,(EF_1 - EF_0)\min(t, 10) && \text{(Exogenous forcing)}
\end{aligned}$$

[41] The deterministic evolution of the variance is in a Markov sense. The agent knows the variance one decade ahead since it is entirely a function of time t variables, but she does not know it two decades ahead since T_{t+1} is unknown.

Table A-1 reports the values of the model parameters. The feedback mean is calibrated so the implied climate sensitivity of 3°C matches the DICE parameterization. We calibrate the variance of the feedback term using the value in Roe and Baker (2007). The variance of the decadal temperature shock is obtained from Kelly and Kolstad (1999b).

Table A-1: The parameters of the dynamic stochastic DICE-2007 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
b_2	0.0028	Damage coefficient
b_3	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 in the entangled preferences case
ϕ_{11}	0.811	Carbon transfer coefficient for atmosphere to atmosphere
ϕ_{12}	0.189	Carbon transfer coefficient for atmosphere to upper ocean
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

Continued on next page

Table A-1 – continued from previous page

Parameter	Value	Description
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)
λ_0	1.2	Reference system's climate sensitivity (°C/[W/m^2])
σ_T^2	0.11	Temperature shock variance
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 (°C)
T_0^{ocean}	0.0068	Year 2005 ocean temperature (°C)
μ_0	0.6	Year 2005 feedback prior mean
Σ_0	0.13^2	Year 2005 feedback prior variance

We make two changes of variables to the model to reduce the computational burden. These changes do not alter the DICE model but only how it is represented in the computer code. In particular, we express both capital, K_t, and consumption, C_t, in terms of effective labor and technology:

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

We keep utility in standard (not effective) terms so that the agent's utility function is

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

The resulting problem yields the exact same solutions as the GAMS version of the DICE model, once modified to use a Markov representation of the forcing equation.

Our dynamic programming version of the DICE model is thus:

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

subject to transitions:

$$
\begin{aligned}
k_{t+1} &= \frac{1}{A_{t+1}^{1/(1-\kappa)}\, L_{t+1}} \left[(1-\delta_k)^{10}\, A_t^{1/(1-\kappa)}\, L_t\, k_t + 10\, \left((1-\psi_t \alpha_t^{a_2}) Y_t^n - A_t^{1/(1-\kappa)}\, L_t\, c_t\right)\right], \\
\begin{bmatrix} M_{t+1}^{atm} \\ M_{t+1}^{up} \\ M_{t+1}^{lo} \end{bmatrix} &= \begin{bmatrix} \phi_{11} & \phi_{21} & 0 \\ \phi_{12} & \phi_{22} & \phi_{32} \\ 0 & \phi_{23} & \phi_{33} \end{bmatrix} \begin{bmatrix} M_t^{atm} \\ M_t^{up} \\ M_t^{lo} \end{bmatrix} + \begin{bmatrix} e_t \\ 0 \\ 0 \end{bmatrix}, \\
T_{t+1}^s &= T_t^s + C_1 \left[F_{t+1}(M_{t+1}^{atm}) - f_{2x} \frac{1-\Delta}{\lambda_0} T_t^s + C_3\,(T_t^o - T_t^s)\right] + \epsilon_t, \\
T_{t+1}^o &= C_4\, T_t^s + (1-C_4)\, T_t^o, \\
\mu_{t+1} &= \frac{\Sigma_t\, \gamma_t\, H_t + \sigma_T^2\, \mu_t}{\Sigma_t\, \gamma_t^2 + \sigma_T^2}, \\
\Sigma_{t+1} &= \frac{\Sigma_t \sigma_T^2}{\Sigma_t\, \gamma_t^2 + \sigma_T^2},
\end{aligned}
$$

where

$$
\begin{aligned}
\gamma_t &= \frac{C_1\, f_{2x} T_t^s}{\lambda_0}, \\
H_t &= T_{t+1}^s - \left(T_t^s + C_1 \left[F_{t+1}(M_{t+1}^{atm}) - f_{2x} \frac{1}{\lambda_0} T_t^s + C_3\,(T_t^o - T_t^s)\right)\right].
\end{aligned}
$$

Finally, we constrain industrial emissions to be nonnegative and we impose the resource constraint:

$$
\begin{aligned}
&\alpha_t \le 1, \\
&A_t^{1/(1-\kappa)}\, L_t\, c_t + (\psi_t \alpha_t^{a_2})\, Y_t^n \le Y_t^n.
\end{aligned}
$$

We take the double expectation over the feedback and temperature shock distributions by using Gauss-Hermite quadrature, with 7 unique quadrature points for the feedback distribution and 7 unique quadrature points for the temperature shock. This results in 49 total quadrature points covering the joint distribution. Our results are not sensitive to varying the number of quadrature points to either 25 or 81.

We solve the model by using value function iteration with a finite horizon. We approximate the value function using the Smolyak algorithm. To implement the Smolyak algorithm, we use slightly altered versions of the software provided by Judd et al. (2014b). We set the

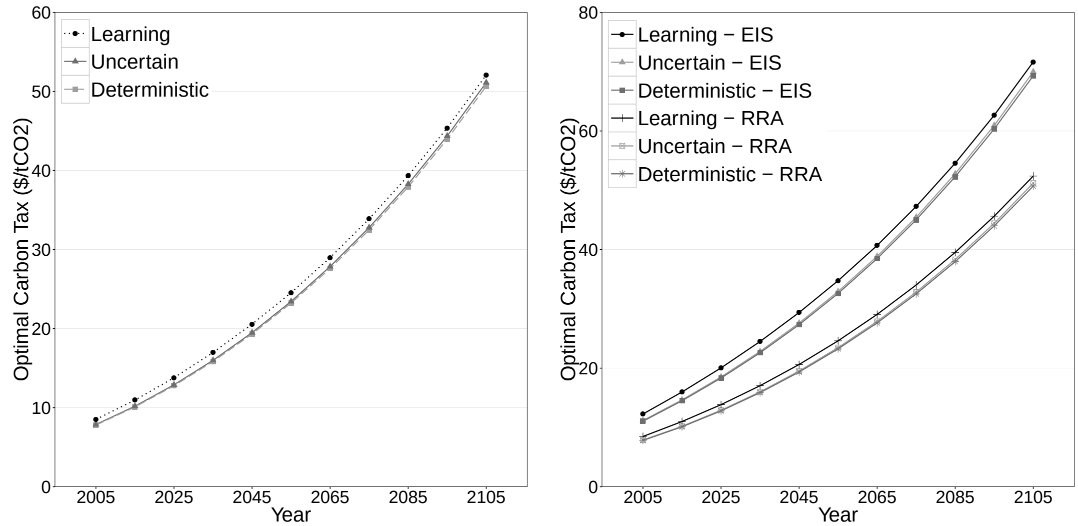

Figure A-2: The optimal emission tax with standard preferences (left) and with disentangled preferences (right). EIS indicates that the reciprocal of the elasticity of intertemporal substition has been lowered to 3/2, and RRA indicates that relative risk aversion has been raised to 3. Random variables happen to take on their mean values in every period.

horizon at 2555 and use a terminal value function where the policymaker has her initial 2005 beliefs about the uncertain parameter and all exogenously changing variables are held constant at their year 2555 values. In order to match the DICE model, we use a 10-year timestep, but we recognize that an annual timestep could be important for a more thorough exploration of learning. Figure A-2 depicts the optimal emission tax trajectory in a deterministic version of our setting, when climate sensitivity is uncertain but the policymaker does not learn from observations, and when the policymaker anticipates learning from observations.

Finally, we extend the theoretical analysis of Section 3 to account for the additional states in DICE. The marginal cost of time t emissions is

$$-\beta E_t\left[\frac{\partial V_{t+1}}{\partial M_{t+1}^{atm}}\frac{\partial M_{t+1}^{atm}}{\partial e_t}+\frac{\partial V_{t+1}}{\partial T_{t+1}^{s}}\frac{\partial T_{t+1}^{s}}{\partial M_{t+1}^{atm}}\frac{\partial M_{t+1}^{atm}}{\partial e_t}\right].$$

Passing the expectation operator through, this becomes:

$$-\beta\left\{\underbrace{E_t\left[\frac{\partial V_{t+1}}{\partial T_{t+1}^{s}}\right]}_{A}\frac{\partial T_{t+1}^{s}}{\partial M_{t+1}^{atm}}\frac{\partial M_{t+1}^{atm}}{\partial e_t}+\underbrace{E_t\left[\frac{\partial V_{t+1}}{\partial M_{t+1}^{atm}}\right]}_{B}\frac{\partial M_{t+1}^{atm}}{\partial e_t}\right\}.$$

Term A is essentially as analyzed in the main text. Term B can be analyzed in a directly analogous fashion. For the main text's figure, we combine the channels in term A with the corresponding channels in term B. We lose the insurance term because $\partial T_{t+1}^{s}/\partial M_{t+1}^{atm}$ is not uncertain from the perspective of time t: the uncertain parameter Δ does not interact with forcing in the temperature transition. We also lose the active learning channel because neither the transition equation for μ_{t+1} nor the transition equation for Σ_{t+1} depends on time t emissions, once we substitute for T^s_{t+1} in H_t. The active learning channel and the insurance channel both reappear after a longer interval lapses, so that they help to determine the adjustment for future uncertainty.

The second-order Taylor expansions of terms A and B appear to be adequate. The maximum relative error between the sum of our Taylor approximation terms and the actual optimal carbon tax is 2×10^{-4} for the learning results in the main text. The average relative error is 8×10^{-5}. The maximum and average relative errors for the uncertainty results are 2×10^{-6} and 2×10^{-7}.

Best Practices

Quantifying errors in solutions and providing sufficient information for replication is critical. Here we demonstrate some best practices using the dynamic stochastic version of DICE developed for this paper. We first demonstrate the accuracy of our model. There are several ways to test model accuracy, but we begin by comparing optimal trajectories from deterministic runs of our dynamic programming model to the deterministic DICE-2007 model solved in GAMS.[42] In order to make an apples to apples comparison, we change the GAMS model to make it Markov: we rewrite the forcing equation so that current forcing is a function of current atmospheric CO_2 and not the average of current and next period's atmospheric CO_2.[43]

Figure A-3 displays the optimal trajectories for the two models from 2005 to 2205, focusing on the abatement rate control, the consumption control, the surface temperature state, and the atmospheric CO_2 state. In these plots, we solve the dynamic programming model on the larger domain that we use when the climate sensitivity is uncertain, and we set $\mu = 3$ for the Smolyak algorithm.[44] Solving the problem on a tighter domain, or omitting the unnecessary belief states for this setting, would indeed lead to greater accuracy but would not provide a fair assessment of the accuracy of the solutions used to analyze uncertainty about

[42] Cai et al. (2013) demonstrate accuracy of their dynamic programming solution by comparing deterministic runs of their dynamic programming model to solutions of their model solved using more accurate optimal control techniques.

[43] Replicating known results is not always a good check. Many recursive IAMs use an annual timestep, whereas the DICE model uses a ten-year timestep. Optimal policies should vary between these settings. Further, finding only small errors in simulated trajectories in one parameterization of the model (such as the standard DICE parameterization) does not guarantee that simulations traveling to other areas of the state space will be as accurate.

[44] The model solves in 15-20 minutes on a 28 core machine. $\mu = 3$ implies 9 unique collocation points for each dimension, although there are not 9^8 points on the resulting grid since it is not constructed from a full tensor product.

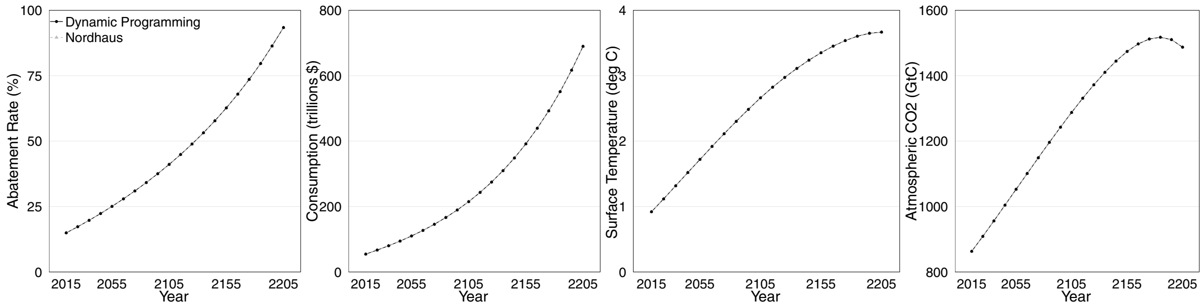

Figure A-3: Optimal trajectories for the GAMS DICE-2007 model (triangles) and our dynamic programming version of DICE-2007 (circles) under the standard DICE parameterization.

Table A-2: Relative errors between GAMS solution and the dynamic programming solution to the DICE-2007 model.

	Abatement Rate	Consumption	Temperature	CO_2
Maximum Relative Error	3.9×10^{-3}	7.2×10^{-4}	3.6×10^{-4}	3.0×10^{-4}
Average Relative Error	1.1×10^{-3}	2.8×10^{-4}	2.2×10^{-4}	1.9×10^{-4}

climate sensitivity.[45]

All plotted state and control trajectories are extremely close, and visually indistinguishable. Table A-2 displays the maximum and average relative errors between the GAMS and dynamic programming trajectories. The abatement rate displays the greatest error out of all the trajectories but is still very accurate and diverges from the GAMS solution by one-tenth of a percent on average. The maximum and average errors for the other trajectories are extremely small and only a few hundredths of a percent.

In addition to checking whether the model has acceptable error in optimal trajectories, we can check for internal consistency in the model.[46] In many macroeconomic models this is done by analyzing the Euler equation residuals. Santos (2000) notes that under certain conditions the size of the Euler residuals is the same order of magnitude as the approximation error in the policy function. However, the DICE model does not satisfy these conditions due to the concavity of forcing and to its use, in some cases, of a finite horizon. Instead we can study residuals of other equations that must be satisfied along an optimal trajectory. These do not have a direct mapping into errors in policy or welfare, but nonetheless can give us

[45] Another way to validate the model and code accuracy is to simulate our model forward using the optimal trajectories from the Markov version of the GAMS DICE-2007 model. Doing this, we find that the state trajectories are nearly identical.

[46] One should always examine the value function derivatives to ensure that they are sensible, and one should always ensure that the solution is not particularly sensitive to the domain or degree of approximation.

Table A-3: First-order condition errors of the learning model in log10 units.

	Abatement Rate	Consumption
Maximum Error	-4.75	-6.70
Average Error	-5.84	-7.31

some insight into the accuracy of our approximation. Judd et al. (2014a) and Fernández-Villaverde et al. (2016) give some suggestions regarding residuals to test, such as Bellman errors, first-order condition errors, or a χ^2 accuracy test.

Here we demonstrate testing first-order condition errors. The first-order conditions determine optimal policy, so ensuring that this error is small along simulated trajectories is critical for getting the correct policy implications of changes in risk and uncertainty. Following the convention for residual analysis, we calculate first-order condition errors by rearranging the first-order conditions so that all terms are on the left-hand side of the equation and 1 is on the other side, subtracting 1 from the left-hand side, and taking the base 10 logarithm of the absolute value. If there were no approximation error, then taking the base 10 logarithm would produce negative infinity, but this does not happen in practice due to errors in the approximation of the value function and due to truncation and rounding during simulation. Table A-3 reports the error in the first-order conditions of our learning model's simulated trajectory in log10 units. Numbers that are larger in magnitude imply smaller errors in the first-order conditions. Along the optimal trajectories, the errors in the first-order conditions are small. The average abatement error is -5.84, and all are smaller than -4.75. Consumption errors are even smaller (-7.31) on average, and all consumption errors are less than -6.70.

Finally, we report the domain of approximation in Table A-4. The domains for both temperature states are matched since a tighter ocean state domain will always result in the ocean state transitioning outside the domain because next periods ocean temperature is a convex combination of current surface and ocean temperature. We select the domain for the two ocean CO_2 states by using the transition equations to infer the tightest bounds that would keep the next period's ocean CO_2 states within the domain.[47] Results do not change if we select a tighter domain for the ocean CO_2 states and allow the continuation value to be evaluated outside the domain during the approximation step. The lower bound for the atmospheric CO_2 state was selected so that, conditional on how we construct the domain for the ocean CO_2 states, the upper ocean CO_2 state does not immediately jump outside the domain during simulation. The variance of beliefs is naturally bounded between the initial belief and zero. The effective capital domain is selected based on the tightest domain that would allow for convergence of the learning model. We select the domain for the mean of the feedback distribution to be accurate under the test against the GAMS version of DICE

[47] Due to numerical error, some grid points do transition outside the domain, but only by a very small amount.

Table A-4: Domain bounds for the approximation space.

	k_t	T_t^s	T_t^o	M_t^{atm}	M_t^{up}	M_t^{lo}	μ_t	Σ_t
Lower Bound	1.7	0	0	580	1229	18204	0.4	0
Upper Bound	6	10.6	10.6	1700	2310	47064	0.8	0.13^2

and to take advantage of the fact that we only simulate the model with all random variables fixed at their mean values. Note that, during the value function approximation steps, the continuation value is evaluated outside the domain at a number of the quadrature points over next period's feedback mean. Since our results are not sensitive to the choice of the number of quadrature points (and thus are not sensitive to the number of mean quadrature points that are evaluated outside or inside the domain during approximation), we take our results to be sufficiently accurate. A wider domain will be required if a user wants to explore simulations with noise shocks or with a true value for the feedback parameter that is different than 0.6.