Tipping Points and Ambiguity in the Economics of Climate Change*

Derek M. Lemoine[†] & Christian Traeger[§]

[†] Department of Economics, University of Arizona
McClelland Hall 401, 1130 E Helen St, Tucson, AZ, 85721-0108, USA
dlemoine@email.arizona.edu

[§] Department of Agricultural & Resource Economics, University of California, Berkeley
207 Giannini Hall #3310, Berkeley, CA 94720-3310, USA
traeger@berkeley.edu

CUDARE Working Paper 1111

This Version: October 13, 2011, First Version: October 2010

We model optimal policy when the probability of a tipping point, the welfare change due to a tipping point, and knowledge about a tipping point's trigger all depend on the chosen policy path. Analytic results show that optimal policy primarily depends on the total welfare change incurred by a tipping point and on policy's ability to affect a tipping point's probability. Simulations with a numerical climate-economy model show that tipping points increase the optimal near-term carbon tax by up to XX% and that the resulting policy path lowers peak warming by XX°C. Different types of possible tipping points can have qualitatively different effects on policy, demonstrating the importance of explicitly modeling the effects of tipping points on system dynamics. Analytically, aversion to ambiguity in the threshold's distribution can amplify or dampen the effect of tipping points on optimal policy, but in our numerical model, ambiguity aversion always slightly increases the near-term optimal carbon tax.

JEL Codes: Q54, D81, Q48, Q00, D61, D62, D90, C63

Keywords: tipping point, threshold, regime shift, ambiguity, climate, uncertainty, integrated assessment, dynamic programming, cost of carbon, carbon tax

*We are grateful for comments from Larry Karp, Ujjayant Chakravorty, and participants at a number of seminars, including the 2011 NBER Environmental and Energy Economics Summer Institute.

1 Introduction

The threat of tipping points in the climate system often serves as an argument for ambitious greenhouse gas emission reductions, but the economic models used to estimate the optimal carbon tax have long assumed that the climate system changes only smoothly (Nordhaus, 1992; Hall and Behl, 2006). These dynamic models couple an economic growth model to a climate model in order to assess how greenhouse gas externalities should influence consumption, investment, and emissions. National and international policy in fact now relies on these models' estimates of the social cost of carbon dioxide emissions (Greenstone et al., 2011). However, by assuming smooth changes in the climate, these models' estimates are probably biased low. Scientists have grown increasingly concerned about the possibility of abrupt changes due to crossing thresholds in the climate system (Alley et al., 2003; Overpeck and Cole, 2006; Lenton et al., 2008), but absent formal analysis, tipping point arguments can support almost any adjustment to the models' estimates. We demonstrate how explicitly modeling tipping points introduces new terms into the optimal carbon tax. Recognizing that tipping point probabilities are themselves rather uncertain given the current state of scientific knowledge, we also show that it is not a priori clear whether aversion to tipping point ambiguity would raise or lower the optimal carbon tax. We then reformulate a benchmark climate-economy integrated assessment model to include the possibility of tipping points in the climate system. We find that this possibility can raise the optimal carbon tax by over 50% and that different types of tipping points from the scientific literature can have qualitatively different effects on optimal policy. Ambiguity aversion does further raise the optimal carbon tax, but we find that ambiguity attitude is less important for optimal policy than is the type of tipping point considered.

Our analytic and numerical modeling fully endogenizes the probability that a tipping point occurs, the ability to learn about the threshold's location, and the welfare loss from crossing the threshold.[2] Previous work has not endogenized all of these elements in a single model. The macroeconomics and real options literatures have studied optimal policy when monetary policy rules or demand dynamics might shift discontinuously in the future (Guo et al., 2005; Davig and Leeper, 2007). However, these regime shifts were controlled by exogenously fixed transition probabilities independent of the model's control variables. Because our regime shifts are triggered by crossing a threshold in the state space, our policymaker directly controls the probability of crossing a threshold.[3] An endogenous regime shift can induce both hedging and averting behavior while an exogenous regime shift can induce only hedging. Further, we endogenize learning about the threshold's location by using a dynamic programming setting. Keller et al. (2004) numerically modeled a threshold that increases the damages caused by climate change, but this threshold's

[2]We use "tipping point" to refer to an irreversible change in system dynamics that occurs upon crossing a threshold. Other notions of tipping points in the economics literature include shifts between multiple equlibria due to small changes in parameters (e.g., Schelling, 1971; Azariadis and Drazen, 1990; Katz and Shapiro, 1994; Ellison and Fudenberg, 2003) and shifts between optimal policy paths due to small changes in initial conditions (e.g., Brock and Starrett, 2003; Mäler et al., 2003; Wagener, 2003).

[3]The natural resources literature has considered the possibility of an endogenous, stochastic threshold (Heal, 1984; Brozović and Schlenker, 2011; Polasky et al., 2011; de Zeeuw and Zemel, 2011). We consider the case of a fixed threshold with optimal learning about its location. Our policymaker can choose to eliminate the chance of a tipping point by keeping the state variables in a region already found to be safe.

distribution could not be updated over time. In contrast, our tipping points affect the climate system itself, and our policymaker optimally updates the threshold's distribution by observing whether it was crossed in the past. Finally, previous work considered "climate catastrophes" that impose an exogenously-defined flow of disutility upon crossing some threshold (Heal, 1984; Tsur and Zemel, 1996; Gjerde et al., 1999; Nævdal, 2006; Nævdal and Oppenheimer, 2007). Our tipping points do not directly affect utility but instead affect the dynamics of the climate system. The welfare consequences of our tipping points depend on the state variables produced by pre-threshold policy and on how post-threshold policy adjusts for the altered dynamics.

Because the probability of incurring a tipping point is far less understood than, for instance, interannual variability in temperature, we might expect optimal decisions to consider the uncertainty about the probabilities governing the threshold outcome. Distinguishing between degrees of uncertainty reaches back to Keynes (1921: Chapter VI) and Ellsberg (1961), who added a confidence weight to probabilities in order to describe uncertainty more comprehensively. Recent theories of ambiguity recognize that many decision-makers evaluate lotteries differently depending on how sure they are of the probabilities governing the payoffs. We extend the Klibanoff et al. (2005, 2009) model of smooth ambiguity to a new setting with poorly understood thresholds. We demonstrate that ambiguity aversion in fact has an ambiguous effect on the optimal carbon tax. To determine the sign of ambiguity aversion's effect, we also extend the dynamic integrated assessment model to use the smooth ambiguity model. Previous work combining ambiguity with climate-economy models either used a static setting (Lange and Treich, 2008) or only evaluated exogenously-given policy paths (Millner et al., 2010). Our dynamic model optimally selects policy paths and updates the ambiguous distribution in response to new information about the threshold's location. We find that ambiguity aversion does in fact increase the optimal carbon tax, but only by a relatively small amount.

We next describe the approach to modeling endogenous tipping points. We then analyze the determinants of the marginal benefit from a change in policy as well as how ambiguity aversion interacts with tipping point possibilities. This marginal benefit will define the optimal carbon tax in our climate application. In Section 4, we reformulate a benchmark climate-economy integrated assessment model by giving it a recursive structure, generalizing its welfare evaluation, and introducing the endogenous possibility of abrupt shifts in system dynamics. In Section 5, we describe the implications of different types of tipping points for the optimal carbon tax path. We conclude in Section 6 with a discussion of the implications for climate policy and economics. The appendix completes the description of the numerical model.

2 Modeling tipping points

We demonstrate how to model endogenous tipping points with endogenous learning about the threshold and an endogenous welfare impact from a tipping point's occurrence. Tipping points are here irreversible shifts in system dynamics that occur upon crossing a threshold in the state space. The probability of a tipping point occurring is endogenous because it depends on the evolution of the state variables, which is in turn determined by the chosen policy path. Crossing the threshold shifts the world from the "pre-threshold" regime to a "post-threshold" regime with permanently altered system dynamics. Because the policymaker lacks access to a structural model of the processes that drive tipping points, these tipping points reveal themselves by changing the reduced-form model used to represent the system's conventional behavior. Once a threshold is crossed, the decision-maker recognizes the event and adjusts policy accordingly. These optimal policy adjustments combine with the post-threshold system dynamics to endogenously determine the welfare change from crossing a threshold. Prior to crossing a threshold, the policymaker knows how he would adjust to the threshold crossing and optimizes that risk over time. Finally, in each period the policymaker updates his distribution for the threshold's location by observing whether the tipping point has occurred or not. The more of the state space that she explores without crossing the threshold, the more likely the threshold is to be in the remaining portion of the state space. As long as the policymaker would never choose to explore some portion of the state space even in the absence of tipping point possibilities, our model is equivalent to one in which there is some probability that the threshold does not exist.

The policymaker solves an infinite-horizon dynamic optimization problem. The optimal policy program at time t depends on the vector S_t of state variables and on the $n \times n$ diagonal matrix ϵ_t of random shocks. Let $V_\psi(S_t, \epsilon_t)$ be the value of the optimal policy program at time t, which can be expressed recursively as a function of the value of the optimal policy program at time $t+1$. The threshold is a hyperplane that divides the state space into the pre-threshold and post-threshold regions. Each region has its own dynamics (i.e., its own regime), with the tipping point's effect being to change dynamics from the pre-threshold regime to the post-threshold regime. The parameter ψ indicates whether V is the value function for the pre-threshold regime ($\psi = 0$) or for the post-threshold regime ($\psi = 1$). Once the state variables cross into the post-threshold region, the threshold becomes irrelevant for future policy and the entire state space takes on the post-threshold dynamics. Because dynamics can change in only one direction, the tipping point is irreversible.

Consider first the post-threshold regime. Once in that world, the pre-threshold regime is no longer relevant because we assume that the effects of crossing a threshold are irreversible: one cannot undo a change in, for instance, climate dynamics simply by returning the state variables to their former values. In this post-threshold world, the Bellman equation defining the value function is:

$$\begin{aligned} V_{\psi=1}\left(S_t, \epsilon_t\right) = \max_{x_t} & \left\{ u\left(x_t, \epsilon_t\, S_t\right) + \beta \int V_{\psi=1}\left(S_{t+1}, \epsilon_{t+1}\right) d\,\mathbb{P} \right\} \\ \text{s.t. } & S_{t+1} = g_{\psi=1}\left(x_t, S_t\right) \\ & f(x_t) \leq Q(\epsilon_t\, S_t) \ . \end{aligned} \tag{1}$$

Here, x_t is a vector of time t control variables, β is the per-period discount factor, and $u(\cdot)$ gives the immediate utility from the time t controls, shocks, and state variables. The function $g_{\psi=1}(\cdot)$ gives the post-threshold transition equation defining the next period's state variables S_{t+1}. The term $Q(\epsilon\, S)$ is the resource constraint as a function of the state variables and the random shocks. The diagonal elements in ϵ_t have joint probability measure $\mathbb{P}$ and are independently and identically distributed over time. In addition, each diagonal element has a mean of 1, giving $E_t[\epsilon_{t+1}\, S_{t+1}] = S_{t+1}$. The noise captured by ϵ produces standard risk, which we will contrast with the ambiguity posed by an unknown threshold. The function $V_{\psi=1}$ will determine the welfare change from crossing a threshold, which in turn depends on the state variables at the time of crossing and on the policies chosen in equation (1).

Prior to crossing a threshold, the value of the optimal policy program is given by the pre-threshold value function:

$$V_{\psi=0}\left(S_t, \epsilon_t\right) = \max_{x_t} \left\{ u\left(x_t, \epsilon_t\, S_t\right) + \beta \int \left[h\left(S_t, S_{t+1}\right)\; V_{\psi=1}\left(S_{t+1}, \epsilon_{t+1}\right) \right.\right.$$
$$\left.\left. + \left[1 - h\left(S_t, S_{t+1}\right)\right]\; V_{\psi=0}\left(S_{t+1}, \epsilon_{t+1}\right) \right] d\,\mathbb{P} \right\}$$
$$\text{s.t. } S_{t+1} = g_{\psi=0}\left(x_t, S_t\right)$$
$$f(x_t) \leq Q(\epsilon_t\, S_t)\ .$$

The pre-threshold value function captures learning and the endogenous threshold possibility. When a threshold might be crossed before the next period, the policymaker must consider both possible value functions when determining the expected value for the next period. The function $h\left(S_t, S_{t+1}\right)$ is the hazard rate. This term gives the probability of crossing a threshold between times t and $t+1$. The hazard rate depends on S_t because this vector (along with the value of ψ) includes all information available at time t about the threshold's location. S_t determines both the region of the state space already shown to be free of the threshold and the probability of the threshold being at any given location in the remainder of the state space. The hazard rate also depends on S_{t+1} because the next period's state variables determine the set of threshold locations that would trigger a tipping point before the next period: greater jumps into an unexplored region of the state space carry a greater chance of crossing over the threshold. The hazard rate's dependence on S_t defines our ability to learn about the threshold's location over time, and the dependence on S_{t+1} endogenizes a tipping point's occurrence by making it depend on the chosen policies. The probability of crossing a threshold does not directly depend on the transitory shock to the state variables: the random noise ϵ_{t+1} affects how the state variables in S_{t+1} enter into the time $t+1$ immediate utility and resource constraint, but it does not affect the chance of crossing a threshold before time $t+1$. We can solve for $V_{\psi=0}$ using standard methods as long as we have already solved for $V_{\psi=1}$. In this case, $V_{\psi=1}$ is a known function, acting like a scrap value that depends on the state variables and is incurred with probability h.

Finally, in an original application of ambiguity aversion, we also generalize the welfare evaluation to recognize that we often know little about the threshold's location. By assumption, we have no records of having crossed the threshold before, and abrupt changes in system dynamics may not lend themselves to accurate forecasting or simulation. In many cases, these abrupt changes emerge from complex systems with multiple, interacting components having an incompletely constrained structure. Tipping point possibilities therefore demonstrate ambiguity because the policymaker has low confidence in the distribution assigned to the threshold. In contrast, the random shocks described by $\mathbb{P}$ are well understood, either because their behavior is constrained by a long time series of observations or because the stochasticity is due to well-understood mechanisms. The random shocks in ϵ therefore pose standard risk because the policymaker has high confidence in the distribution placed on them. An ambiguity-averse decision-maker would evaluate the tipping point uncertainty differently from the uncertainty about the realization of the random shock. We therefore generalize the pre-threshold welfare evaluation so that the policymaker is not constrained to be ambiguity-neutral:

$$
\begin{aligned}
V_{\psi=0}(S_t,\epsilon_t) = \max_{x_t} \Bigg\{ u(x_t, \epsilon_t S_t) + \beta \int f_{amb}^{-1} \Big[& h(S_t,S_{t+1}) \;\; f_{amb}\left[V_{\psi=1}(S_{t+1},\epsilon_{t+1})\right] \\
& + \left[1-h(S_t,S_{t+1})\right] \;\; f_{amb}\left[V_{\psi=0}(S_{t+1},\epsilon_{t+1})\right] \Big] d\,\mathbb{P} \Bigg\} \qquad (2) \\
\text{s.t. } \; S_{t+1} &= g_{\psi=0}(x_t,S_t) \\
f(x_t) &\leq Q(\epsilon_t S_t) \; .
\end{aligned}
$$

The concave function f_{amb} captures smooth ambiguity aversion (Klibanoff et al., 2005, 2009), or intertemporal risk aversion to subjective uncertainty (Traeger, 2010). When f_{amb} is linear, equation (2) reduces to the ambiguity-neutral form. Behavioral evidence supports a concave f_{amb} as a descriptive model of many decision-makers (Camerer and Weber, 1992). Further, when the von Neumann-Morgenstern axioms are extended to recognize that a decision-maker can have differing confidence in different distributions, concave f_{amb} is consistent with normatively attractive preferences (Traeger, 2010).

3 The effects of tipping points on optimal policy

We now identify the channels by which tipping points affect optimal policy. We show how tipping point possibilities add two new terms to the marginal welfare impact of changing a control. In the climate application, tipping point possibilities therefore add two new terms to the social cost of carbon (optimal carbon tax) omitted by previous work assuming a smooth climate system.

The optimal level of the time t control k_t depends on the change in welfare produced by a marginal change in k_t. For ease of exposition, we assume that the only time $t+1$ state variable affected by this change is the single element of S_{t+1} that we denote by j_{t+1}. At an interior optimum, the marginal change in welfare due to an increase in k_t equals the product of the marginal change in the resource function f and the shadow value of the resource constraint Q. When analyzing the social cost of carbon in the climate change application (Section 4), k_t is time t emissions. These emissions will affect two state variables: the time $t+1$ CO_2 stock and time $t+1$ temperature. The benefit from a marginal reduction in time t emissions defines the social cost of carbon and, along the optimal policy path, equals the cost of allocating additional resources to further reduce emissions.

We analyze the right-hand side of equation (2) at the optimal controls (denoted by *). Suppressing all arguments independent of k_t, the value of the optimal policy program is:

$$
u(k_t^*) + \beta_t \int \underbrace{f_{amb}^{-1}\Big[\left[1-h(j_{t+1}(k_t^*))\right] f_{amb}\left[V_{\psi=0}(j_{t+1}(k_t^*))\right] + h(j_{t+1}(k_t^*)) f_{amb}\left[V_{\psi=1}(j_{t+1}(k_t^*))\right]\Big]}_{V_{eff}(k_t^*)} d\mathbb{P} \; .
$$

The change in value from an additional (marginal) unit of k_t is:

$$\begin{aligned}\frac{\partial u(k_t^*)}{\partial k_t} + \beta \int \Bigg\{ & \left[1 - h(j_{t+1}(k_t^*))\right] \frac{f'_{amb}\left[V_{\psi=0}(j_{t+1}(k_t^*))\right]}{f'_{amb}\left[V_{eff}(k_t^*)\right]} \frac{\partial V_{\psi=0}(j_{t+1}(k_t^*))}{\partial j_{t+1}} \frac{\partial j_{t+1}(k_t^*)}{\partial k_t} \\ & + h(j_{t+1}(k_t^*)) \frac{f'_{amb}\left[V_{\psi=1}(j_{t+1}(k_t^*))\right]}{f'_{amb}\left[V_{eff}(k_t^*)\right]} \frac{\partial V_{\psi=1}(j_{t+1}(k_t^*))}{\partial j_{t+1}} \frac{\partial j_{t+1}(k_t^*)}{\partial k_t} \\ & + \underbrace{\frac{\partial h(j_{t+1}(k_t^*))}{\partial j_{t+1}}}_{\text{(i)}} \underbrace{\frac{\partial j_{t+1}(k_t^*)}{\partial k_{t+1}}}_{\text{(ii)}} \underbrace{\frac{f_{amb}[V_{\psi=1}(j_{t+1}(k_t^*))] - f_{amb}[V_{\psi=0}(j_{t+1}(k_t^*))]}{f'_{amb}\left[V_{eff}(k_t^*)\right]}}_{\text{(iii)}} \Bigg\} d\mathbb{P} ,\end{aligned} \tag{3}$$

where primes ($'$) indicate derivatives. Assume for now that the decision-maker is ambiguity-neutral, making f_{amb} the identity function. In a world without tipping points, the two terms in the first line of equation (3) wholly define the optimal controls because the hazard is 0 for all values of the state variables ($h(S_t, S_{t+1}) = 0 \ \forall S_{t+1}$). This first line gives the marginal effect of control k_t on instantaneous utility and on time $t+1$ welfare under the pre-threshold regime (i.e., on the pre-threshold continuation value). The effect on time $t+1$ welfare is composed of the effect of a change in j_{t+1} on the pre-threshold value function $V_{\psi=0}$ and the effect of a change in k_t on j_{t+1}. Models that assume smooth dynamics include only this first line. The next two lines indicate how tipping point considerations change the marginal welfare effect of control k_t: when tipping points are possible, altering k_t also changes time $t+1$ welfare in the post-threshold world (second line) and changes the probability of entering the post-threshold world (third line). The second line is analogous to the effect on time $t+1$ welfare captured in the first line, except it uses the post-threshold value function $V_{\psi=1}$ and weights by the hazard rate h rather than by $1-h$. The third line gives the effect of a change in the hazard rate on expected time $t+1$ welfare. This expectation is formed by aggregating the pre- and post-threshold value functions using the probability that each obtains at time $t+1$, giving $(1-h)V_{\psi=0} + hV_{\psi=1}$. The change in this expectation due to a change in the hazard rate h is multiplied by the change in the hazard rate due to a marginal change in k_t. We call the net effect the "hazard effect" because it operates through the effect of k_t on the hazard rate rather than through the effect of k_t on the value functions themselves.

The importance of tipping point considerations for optimal policy depends on the magnitude of the second and third lines in equation (3) in comparison to the first line. The second line's contribution is significant when the per-period hazard rate is large and when changing j_{t+1} has quite different effects on welfare in the pre- and post-threshold worlds. In our climate application, the smallness of the annual hazard rate (order 10^{-3}) makes this second line have only a trivial effect on the social cost of carbon and so on the optimal carbon tax. In contrast, the third line (the hazard effect) can have quite a significant effect on the optimal carbon tax. In fact, we expect that the hazard rate will generally be the main channel by which tipping points affect optimal policy. The hazard effect is composed of three terms. Term i gives the effect of a change in the state variable j_{t+1} on the hazard rate, term ii gives the effect of a change in control k_t on the state variable j_{t+1}, and term iii gives the total welfare change from crossing a threshold. In most interesting applications, term iii will be the largest: the first two terms are each marginal effects in continuous functions, but term iii is a total change between two distinct functions. We usually deem thresholds important when we expect their crossing to cause a significant change in welfare. In this case, the primary effect of the threshold on control k_t will usually be channeled through term iii by way of the hazard effect. If a marginal increase in k_t raises the hazard rate (i.e., if increasing k_t makes a tipping point more likely), the sign of the hazard effect is the same as the sign of $V_{\psi=1} - V_{\psi=0}$. For a welfare-increasing tipping point, the hazard effect changes a control's optimal level to make crossing the threshold more likely. For a welfare-decreasing tipping point, the hazard effect leads optimal controls to reduce the probability of crossing the threshold. Therefore, in our climate application, the hazard effect raises the social cost of carbon and so also the optimal carbon tax.

One might expect ambiguity aversion (or increased aversion to tipping point risk) to reinforce the desire to avoid undesirable tipping point outcomes, but we find that this need not be the case. The introduction of ambiguity aversion modifies the terms in equation (3). Consider first the effect of ambigutiy aversion on the first two lines. Note that $V_{eff}(k_t^*)$ is a generalized mean of the pre- and post-threshold value functions and that an ambiguity-averse decision-maker has concave f_{amb}. Therefore, we have $V_{\psi=1} > V_{eff} > V_{\psi=0}$ if a tipping point increases welfare and $V_{\psi=0} > V_{eff} > V_{\psi=1}$ if a tipping point decreases welfare. Consider the case where a tipping point decreases welfare. Here the fractions

$$\frac{f'_{amb}\left[V_{\psi=0}(j_{t+1}(k_t^*))\right]}{f'_{amb}\left[V_{eff}(k_t^*)\right]} < 1 \quad \text{and} \quad \frac{f'_{amb}\left[V_{\psi=1}(j_{t+1}(k_t^*))\right]}{f'_{amb}\left[V_{eff}(k_t^*)\right]} > 1$$

give more weight to the effect of control k_t in the less desirable regime. These fractions decrease the weight given to the pre-threshold effect of k_t (first line in equation (3)) and increase the weight given to the post-threshold effect of k_t (second line in equation (3)). The reverse is true for a tipping point that increases welfare. Thus, ambiguity aversion biases the probability of the less desirable outcome upwards when considering the effect of k_t on future welfare under each regime. We call this reweighting of the time $t+1$ value functions the "pessimism bias" induced by ambiguity aversion. A second-order approximation of the function f_{amb} around V_{eff} together with a linearization of V_{eff} yields the following approximation of the pessimism bias:

$$\frac{1}{2}\frac{\partial j_{t+1}}{\partial k_t}\left.\frac{-f''_{amb}}{f'_{amb}}\right|_{V_{eff}(k_t^*)} h(1-h)\big(V_{\psi=0} - V_{\psi=1}\big)\left(\frac{\partial V_{\psi=0}}{\partial j_{t+1}} - \frac{\partial V_{\psi=1}}{\partial j_{t+1}}\right) . \tag{4}$$

This pessimism bias adds to the ambiguity-neutral contribution $\frac{\partial j_{t+1}}{\partial k_t}\left[(1-h)\frac{\partial V_{\psi=0}}{\partial j_{t+1}} + h\frac{\partial V_{\psi=1}}{\partial j_{t+1}}\right]$ in the first two lines of equation (3). Pessimism bias amplifies the tipping point's effect on optimal policy for a welfare-decreasing tipping point and dampens its effect for a welfare-increasing tipping point.[4] Because pessimism bias increases the weight of the less desirable outcome, it counteracts the influence of considering tipping point possibilities when the less desirable outcome is the world in which a tipping point has not occurred. The magnitude of the pessimism bias is proportional to the measure of absolute ambiguity aversion $-f''_{amb}/f'_{amb}|_{V_{eff}(k_t^*)}$ and to $h(1-h)$. The measure

[4]Note that for an ambiguity-neutral decision-maker, considering a tipping point possibility changes the first two lines of equation (3) by $h\frac{\partial j_{t+1}}{\partial k_t}\left(\frac{\partial V_{\psi=0}}{\partial j_{t+1}} - \frac{\partial V_{\psi=1}}{\partial j_{t+1}}\right)$. The pessimism bias has the same sign when $V_{\psi=o} - V_{\psi=1} > 0$.

of absolute ambiguity aversion is 0 for an ambiguity-neutral decision-maker and is positive for an ambiguity-averse decision-maker. The term $h(1-h)$ reaches its maximum when $h = 0.5$ and reaches its minimum as the hazard rate goes to 1 or 0. Therefore, as the next period's tipping point outcome becomes more certain, ambiguity aversion has less of an impact when evaluating the marginal effect of k_t on welfare under each regime. In our climate application with welfare-decreasing tipping points, pessimism bias increases the optimal carbon tax but only by a trivial amount. Because we model tipping points as unlikely to occur in a given year, pessimism bias is of small numerical importance for the optimal carbon tax.

We have seen that the primary effect of tipping points is often channeled through the hazard effect in the third line of equation (3), and we now see that the same is often true of ambiguity aversion. However, ambiguity aversion has an ambiguous impact on the magnitude of the hazard effect. We assume $f_{amb}[V_{\psi=1}(j_{t+1}(k_t^*))] = V_{\psi=1}(j_{t+1}(k_t^*))$ and $f_{amb}[V_{\psi=0}(j_{t+1}(k_t^*))] = V_{\psi=0}(j_{t+1}(k_t^*))$, which is without loss of generality because the function f_{amb} is only determined up to affine transformations in a preference representation. The denominator $f'_{amb}\big[V_{eff}(k_t^*)\big]$ in term iii in equation (3) then controls how ambiguity aversion changes the hazard effect. Ambiguity aversion implies that f_{amb} is concave. Thus, ambiguity aversion increases the hazard effect most if $V_{eff}(k_t^*)$ is close to the value in the preferred regime. In other words, for welfare-decreasing tipping points, ambiguity aversion increases the hazard effect most if $V_{eff}(k_t^*)$ is close to $V_{\psi=0}(j_{t+1}(k_t^*))$ and it decreases the hazard effect most if $V_{eff}(k_t^*)$ is close to $V_{\psi=1}(j_{t+1}(k_t^*))$.[5]

It might be surprising that ambiguity aversion (or greater risk aversion under less confidently known uncertainty) can possibly decrease the importance of a policy variable's effect on the hazard rate. For instance, in the climate application, ambiguity aversion could conceivably decrease the social cost of carbon by decreasing the hazard effect, meaning that an ambiguity-averse policymaker could prefer a lower carbon tax than would an ambiguity-neutral policymaker. The intuition for the possible dampening of the hazard effect relates to a finding in the standard risk setting that a risk-averse agent is not always willing to pay more for a risk reduction. Imagine that increasing k_t increases the hazard rate for a welfare-decreasing tipping point but also increases future capital. This is the case in the climate context when k_t is emissions. A policymaker then has to give up wealth in order to decrease the hazard of crossing into a lower-welfare post-threshold regime. The value of the unit she has to give up is determined by her expectation about the threshold crossing. If she is very likely to cross the threshold (h is large), the ambiguity-averse policymaker prefers to bring the additional unit of wealth into the post-threshold regime than relinquish it to reduce the hazard rate. If, in contrast, the per-period hazard h is small, then the generalized mean $V_{eff}(k_t^*)$ is close to $V_{\psi=0}(j_{t+1}(k_t^*))$. Effective future welfare is high, and the ambiguity-averse policymaker is happy to spend more to reduce the hazard than would an ambiguity-neutral policymaker.

The contribution of ambiguity aversion to the hazard effect becomes clearer with a second-order approximation of the function f_{amb} around V_{eff}. As with pessimism bias, the change in the hazard

[5]For a welfare-decreasing tipping point, the slope of f_{amb} is always positive and smallest at $V_{\psi=0}(j_{t+1}(k_t^*))$ and largest at $V_{\psi=1}(j_{t+1}(k_t^*))$. Moreover, the intermediate value theorem tells us that somewhere between these two points the slope has to equal $V_{\psi=1}(j_{t+1}(k_t^*)) - V_{\psi=0}(j_{t+1}(k_t^*))$ by our assumption $f_{amb}[V_{\psi=1}(j_{t+1}(k_t^*))] = V_{\psi=1}(j_{t+1}(k_t^*))$ and $f_{amb}[V_{\psi=0}(j_{t+1}(k_t^*))] = V_{\psi=0}(j_{t+1}(k_t^*))$. At this point the ambiguity contribution switches from increasing the hazard effect (for higher values) to decreasing the hazard effect (for lower values).

effect due to ambiguity aversion is proportional to the measure of absolute ambiguity aversion:

$$\frac{f_{amb}[V_{\psi=1}(j_{t+1}(k_t^*))] - f_{amb}[V_{\psi=0}(j_{t+1}(k_t^*))]}{f'_{amb}\left[V_{eff}(k_t^*)\right]} \approx \left[V_{\psi=1} - V_{\psi=0}\right] \left[1 + \frac{-f''_{amb}}{f'_{amb}}\Big|_{V_{eff}(k_t^*)} \delta\right] , \quad (5)$$

where $\delta = V_{eff} - \frac{V_{\psi=1}+V_{\psi=0}}{2} \approx (h - 1/2)\left(V_{\psi=1} - V_{\psi=0}\right)$. Earlier we saw that when a marginal increase in k_t raises the hazard rate, the sign of the hazard effect is the same as the sign of $V_{\psi=1} - V_{\psi=0}$. Ambiguity aversion now raises the hazard effect by a factor proportional to the (positive) measure of absolute ambiguity aversion, to the total welfare change from crossing a threshold, and, approximately, to $h - 1/2$. Therefore, ambiguity aversion amplifies the hazard effect only when the more desirable threshold outcome is also the more probable one. In other words, ambiguity aversion further biases a control's optimal level towards reducing the probability of the less desirable outcome only when the likelihood of the less desirable outcome is already low. In contrast, if the likelihood of a welfare-decreasing tipping point is high, ambiguity aversion makes the decision-maker less inclined to use her wealth to avoid the unfavorable outcome that she deems highly likely anyway. In our climate application, the probability of a tipping point is low and a tipping point decreases welfare, so ambiguity aversion amplifies the hazard effect and raises the optimal carbon tax. Finally, because ambiguity aversion raises the hazard effect by an amount approximately proportional to the total welfare change from switching regimes, the hazard effect will often be the primary channel by which ambiguity aversion affects optimal policy in the presence of tipping points having large welfare consequences.

4 A climate-economy model with endogenous tipping points, learning, and ambiguity

We now consider the effect of climate tipping points on the optimal carbon tax path. We reformulate the benchmark Dynamic Integrated model of Climate and the Economy (DICE) from Nordhaus (2008) as an infinite-horizon dynamic programming problem with an endogenous tipping point in the climate system, optimal learning about the threshold that triggers a tipping point, and a generalized welfare evaluation.[6] DICE is a Ramsey-Cass-Koopmans growth model that has an aggregate world economy interacting with a climate module (Figure 1). Gross economic output (or potential GDP) is determined by an endogenous capital stock, an exogenously growing labor force, and exogenously improving production technology. Gross output produces CO_2 emissions. Non-abated CO_2 emissions accumulate in the atmosphere and ultimately translate into global warming, which causes damage proportional to world output. Cumulative temperature change affects the total output available for allocation by the policymaker. The control variables are abatement and consumption, and residual output not allocated to these two options becomes capital investment. The state variables are capital per effective unit of labor, the stock of CO_2 in the atmosphere, the cumulative change in global mean surface temperature, and, to keep track of exogenously-evolving variables, time t.

[6]The appendix provides the model equations. The standard DICE model is a nonlinear programming problem that produces open-loop policy rules and has constant system dynamics. For previous work using recursive versions of DICE, see Kelly and Kolstad (1999), Leach (2007), and Crost and Traeger (2010).

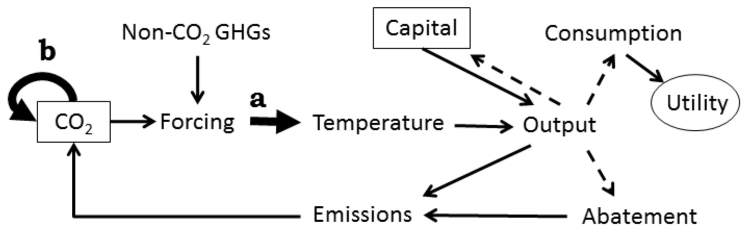

Figure 1: A simplified schematic of the modeled relation between the economy and the climate. Boxes indicate stock variables, and dashed arrows indicate the decision variables of consumption, investment, and abatement. The climate feedback tipping point makes temperature more sensitive to forcing (a), and the CO_2 sink tipping point increases the persistence of CO_2 (b).

A tipping point irreversibly changes the climate system from its conventional representation in DICE to a new regime with altered dynamics. The tipping point occurs upon crossing some unknown temperature threshold. Emission decisions affect whether tipping points occur by determining future temperatures. The decision-maker anticipates how he would choose emissions and consumption in the post-threshold world. The timing, probability, and welfare consequences of a regime switch are endogenous because they depend on the policies chosen before and after the threshold occurs.

To model tipping points, we specialize the recursive structure from Section 2 to DICE. This means that we have one dynamic programming problem for the post-threshold world and another for the pre-threshold world. The pre-threshold world has standard DICE dynamics along with a tipping point possibility and temperature shocks calibrated to the historical record. A tipping point produces the post-threshold world by irreversibly changing the standard dynamics and removing the possibility of crossing additional thresholds. We first solve the post-threshold problem and then use that solution in the pre-threshold problem. We numerically solve each dynamic programming problem for the unknown value function using function iteration. Employing a projection method, we approximate the value functions by Chebychev polynomials and use collocation at the Chebychev nodes in the four-dimensional state space.

We evaluate two tipping points of prominent concern in the climate science literature. The first tipping point increases the feedbacks that amplify global warming (arrow a in Figure 1), and the second increases the atmospheric lifetime of CO_2 (arrow b in Figure 1). The first tipping point therefore increases the degree to which emissions affect temperature, and the second increases the time for which emissions affect the climate. The climate science literature has compiled a number of pathways by which tipping points could abruptly change the strength of feedbacks that determine surface temperature. As one example, warming could mobilize large methane stores locked in permafrost and in ice lattices (clathrates) in the shallow ocean (Hall and Behl, 2006; Archer, 2007; Schaefer et al., 2011). If warming mobilizes these methane stocks, they would cause further warming that could mobilize additional stocks. As another example, if land ice sheets begin to retreat on decadal or centennial timescales, the resulting loss of reflective ice could double the long-term warming predicted using models that hold these ice sheets fixed (Hansen et al., 2008). We represent a feedback tipping point as increasing climate sensitivity, or the equilibrium warming from doubling CO_2, from the standard DICE value of 3°C to 4°C, 5°C, or 6°C. The second tipping point reflects the possibility that carbon sinks weaken beyond the predictions of coupled climate-carbon cycle models. Warming-induced changes in oceans (Le Quéré et al., 2007), soil carbon dynamics (Eglin et al., 2010), and standing biomass (Huntingford et al., 2008) could affect the uptake of CO_2 from the atmosphere. We represent these weakened sinks by decreasing the atmosphere's transfer coefficient by 25%, 50%, or 75%. In every model run, the policymaker faces a single tipping point and knows in advance what its effects would be.

The system passes from the pre-threshold regime ($\psi_t = 0$) into the post-threshold regime ($\psi_{t+1} = 1$) when cumulative temperature change T_{t+1} crosses an unknown threshold $\tilde{T}$. Every temperature between the maximum temperature previously reached and an upper bound $\bar{T}$ has an equal chance of being the threshold, meaning $\tilde{T}$ is uniformly distributed between the historic maximum and $\bar{T}$ (Figure 2a).[7] In our base case model runs, we use $\bar{T} = 4.27$°C so that the year 2005 expected value for the threshold is 2.5°C, and sensitivity analyses vary $\bar{T}$ between 3°C and 9°C.[8] The probability of crossing the threshold between periods t and $t+1$ conditional on not having crossed the threshold by time t is:[9]

$$h(T_t, T_{t+1}) = \max\left\{0, \frac{\min\{T_{t+1}, \bar{T}\} - T_t}{\bar{T} - T_t}\right\} . \tag{6}$$

This expression is the hazard rate for crossing the tipping point. The hazard rate increases as the world reaches higher temperatures without crossing a threshold. This is because the decision-maker learns that the threshold is not below the current temperature, which makes her compress the probability mass into a smaller temperature interval (Figure 2). The hazard rate therefore reflects learning about the threshold's location, and the chosen policies determine each period's hazard rate by determining the next period's temperature.

As noted in discussion of equation (2), the function f_{amb} captures smooth ambiguity aversion (Klibanoff et al., 2005, 2009), or intertemporal risk aversion under subjective uncertainty (Traeger, 2010). We use the term "subjective uncertainty" to describe uncertain outcomes when there is less information available for determining probabilities. This deficiency in probabilistic knowledge applies to climate tipping points, which are less understood than other climate phenomena (Alley et al., 2003; Lenton et al., 2008; Ramanathan and Feng, 2008; Kriegler et al., 2009; Smith et al.,

[7] The optimal policy path in the absence of tipping points reaches a maximum temperature of 3.33°C. Our model with $\bar{T} > 3.33$ is therefore equivalent to one with the uniform distribution's upper bound at 3.33 and probability $(\bar{T} - 3.33)/(\bar{T} - T_t)$ that there is no threshold.

[8] Using $E_{2005}\ \tilde{T} = 2.5$°C is consistent with the political 2°C limits for avoiding dangerous anthropogenic interference. Further, in Smith et al. (2009), 2.5°C is in the upper end of the temperature region that produces significant risk of large-scale discontinuities and is just below the temperatures that produce severe risk.

[9] In DICE-2007, the CO_2 stock increases monotonically until the model reaches a sufficiently high level of abatement. From this point on, the decay rate of CO_2 outweighs the flow of emissions, making the CO_2 stock decrease monotonically. Temperature in DICE follows the same pattern. With increasing temperature, the probability of crossing the threshold is proportional to the difference between the next period's expected temperature and the current temperature. When temperature is decreasing, the probability of crossing the threshold is 0. As long as expected temperature is a quasiconcave function of time, we do not need an additional state variable to keep track of the highest historic temperature.

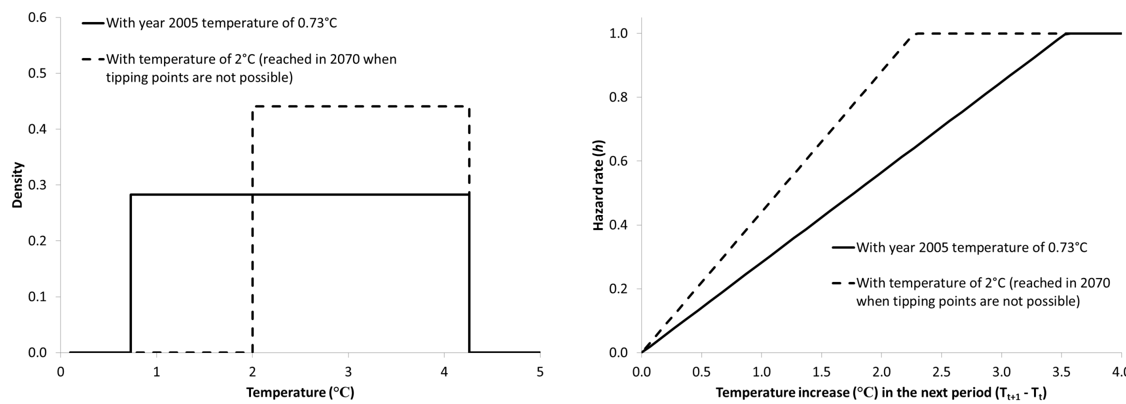

(a) Probability distribution for $\tilde{T}$ with $\bar{T} = 4.27°$C (b) Hazard rate with $\bar{T} = 4.27°$C

Figure 2: As the time t expected temperature increases without crossing a threshold, the probability distribution for the threshold level $\tilde{T}$ places more mass on temperatures yet to be reached. Each additional increase in temperature therefore also produces a greater risk of crossing the threshold.

2009). Sticking to isoelastic preferences, we adopt the power function $f_{amb}(V) = ((1-\eta)V)^{\frac{1-\gamma}{1-\eta}}$. Here γ is a measure of Arrow-Pratt relative risk aversion with respect to subjective or poorly understood uncertainty and η is the constant Arrow-Pratt measure of relative risk aversion used in the standard DICE utility function. The subjective uncertainty contrasts with the standard risk posed by the historically-grounded temperature stochasticity and governed by $\mathbb{P}$. If Arrow-Pratt risk aversion with respect to subjective risk γ coincides with standard risk aversion η, the function f_{amb} is linear and drops out. In that case, the policymaker is ambiguity-neutral and the welfare evaluation is as in DICE. However, when $\gamma > \eta$, the function f_{amb} measures the policymaker's additional aversion to ambiguity (or subjectivity of belief) relative to the better understood risk posed by temperature stochasticity.

5 The optimal carbon tax when facing possible tipping points

We compare several sets of model runs to assess how the optimal carbon tax responds to the type of tipping point considered, to the strength of a tipping point's effect, to prior beliefs about the temperature threshold's location, and to aversion to tipping point ambiguity. The optimal carbon tax (or the social cost of carbon) is determined by the marginal benefit of emission reductions. In Section 3, we analyzed the marginal benefit of a control k_t that affects a state variable j_{t+1}. In determining the marginal benefit of emission reductions, the control is emissions (or, more precisely, abatement), which affects time $t+1$ temperature and CO_2. While the optimal carbon tax in a model without tipping points depends on how emissions affect pre-threshold welfare (the first line in equation (3)), the possibility of tipping points makes the optimal carbon tax also depend on the effect of emissions on post-threshold welfare (the second line in equation (3)) and on the effect of emissions on the probability of crossing a threshold (the third line in equation (3), or the "hazard effect”). While the effect of emissions on post-threshold welfare can be nearly twice as great as the effect on pre-threshold welfare, the smallness of the annual hazard rate along the optimal path (of order 10^{-3} with $\bar{T} = 4.27$) makes the second line's contribution small.

However, additional emissions affect not only future welfare in the the pre- and post-threshold worlds but also the probability that post-threshold world is realized. Additional emissions raise the next period's temperature (term ii in equation (3)). If the next period's temperature is higher than the historic maximum, then the further temperature increase augments the hazard rate (term i). The response of the hazard rate to additional emissions multiplies the cost of the regime switch, which is proportional to the difference between the pre- and post-threshold value functions (term iii). If, on the other hand, the expected temperature is not greater than the historic maximum, then the third line does not contribute to the social cost of carbon because the hazard rate is zero with or without the additional emissions.

The contribution of the third line to the social cost of carbon (the hazard effect) is crucial to understanding the influence of possible tipping points on the optimal carbon tax. The value lost from switching regimes (term iii) is often large. As long as optimal temperatures are weakly increasing (term i is nonzero), the hazard effect can significantly increase the optimal carbon tax relative to a scenario without tipping points (point a in Figure 3). As CO_2 concentrations increase, terms i and iii grow bigger while term ii grows smaller. The changes in terms i and iii generally dominate, increasing the hazard effect over the timespan for which CO_2 and temperature are rising. The increasing hazard effect over time drives an increasingly large wedge between the optimal carbon tax with and without possible tipping points. Eventually, abatement becomes cheap enough that it is worth eliminating threshold risk. At the point where optimal temperature is constant (point b in Figure 3), the social cost of carbon differs between the last unit emitted—for which only the first line in equation 3 contributes—and the next additional unit emitted—for which the first and third lines contribute as temperature would increase once more over the historic level.[10] Once marginal abatement cost falls below the social cost of the last emitted unit (point c in Figure 3), the decision-maker further reduces emissions and so decreases temperature. The social cost of the next and the last ton of carbon emitted merge again, with each only constituted by the first line of equation (3). Finally, once the economy reaches full abatement (point d in Figure 3), the modeled social cost of carbon and marginal abatement cost diverge once more because while the benefit of further emission reductions continues to increase, DICE does not allow negative emission outcomes.

We assess the influence of tipping points on optimal policy by considering the optimal carbon tax path, CO_2 concentration path, and temperature path conditional on not having crossed the threshold (Figure 4). The depicted paths draw the multiplicative temperature shock ϵ at its expected value in each period. We show optimal policy conditional on not having crossed the threshold because we want to understand how optimal policy controls the probability that tipping points occur. Each graph compares the baseline scenario without tipping point awareness to runs with tipping points of various strengths. Each possible tipping point increases the near-term opti-

[10] In a simpler model without a temperature state variable, the emission level that keeps temperature constant changes much more slowly than in the 4-state model we use. Optimal policy in the 3-state model therefore eliminates threshold risk by holding temperature constant for a lengthy period. Eventually abatement cost falls enough to justify decreasing temperature.

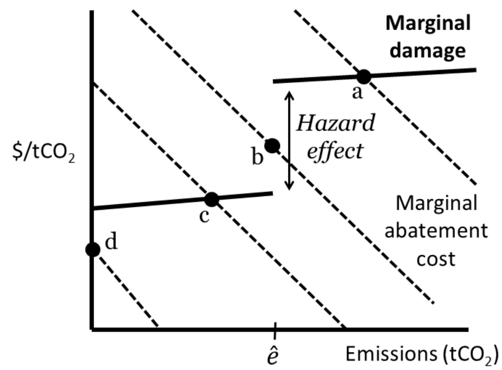

Figure 3: The hazard effect introduces a discontinuity into the marginal damage (i.e., the marginal benefit of abatement) curve at the emission level $\hat{e}$ that keeps temperature constant. Optimal emissions are initially great enough to increase temperature (a). As marginal abatement cost (dashed lines) falls over time, temperature becomes constant (b) before decreasing (c). Eventually, abatement reaches its maximal level (d). Note that marginal damage and the hazard effect both increase over time, and $\hat{e}$ depends on temperature and CO_2.

mal carbon tax by up to 50%, raising its optimal year 2015 value from near \$10/tCO_2 in the base case to near \$14/tCO_2 with the strongest version of either type of tipping point. While tipping point possibilities have only a modest effect on near-term abatement, they can have a large effect on cumulative abatement as the optimal tax rises over time. The tax path without possible tipping points produces a peak CO_2 concentration (temperature) of 637 ppm (3.33°C), reached in the year 2163 (2187). The optimal tax path in the presence of the weaker climate feedback tipping point reduces this peak CO_2 level to 592 ppm (3.03°C), while the higher taxes justified by the stronger climate feedback tipping point further reduce peak CO_2 to 560 ppm (2.81°C). The tax path in the presence of the weaker carbon sink tipping point reduces peak CO_2 only to 617 ppm (3.21°C), while the possibility of the stronger carbon sink tipping point reduces peak CO_2 to 588 ppm (3.02°C). By reducing peak CO_2 and temperature, the decision-maker reduces the cumulative probability of crossing the temperature threshold, and the worse the effects of the anticipated tipping point, the more output the decision-maker devotes to reducing this probability.

Because we explicitly model the effects of tipping points on system dynamics, different types of tipping points can have qualitatively different types of effects on optimal policy. The carbon sink tipping point often has a greater effect on the near-term carbon tax than does the feedback tipping point, yet it also often has a lesser effect on cumulative abatement. The reason is that the total welfare loss (term iii in equation (3)) from crossing a temperature threshold evolves differently under each tipping point. Under the feedback tipping point, high CO_2 concentrations cause much higher damages than in the pre-threshold regime. Further, because optimal post-threshold policy adjusts strongly for the increased sensitivity of temperature to emissions by imposing a much higher carbon tax, optimal CO_2 levels in the post-threshold regime are lower than under the pre-threshold regime. The welfare loss from crossing the feedback tipping point's threshold therefore rises over time as the CO_2 stock also rises. This in turn increases the hazard effect over time, which further

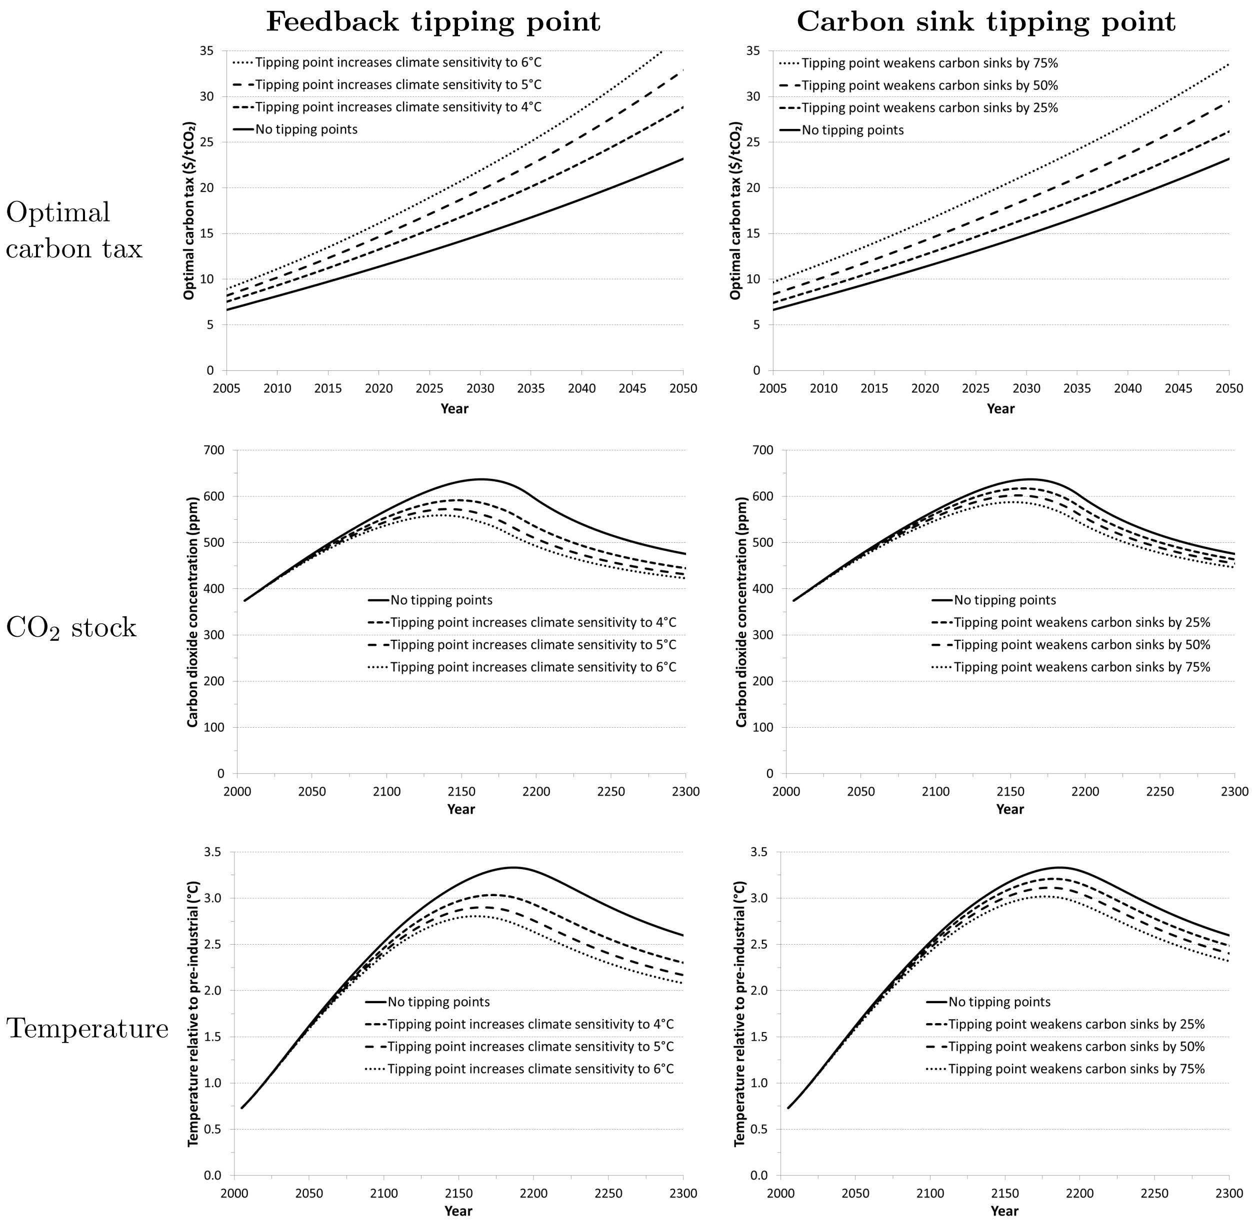

Figure 4: Time paths for the optimal carbon tax (current value), the CO_2 stock, and temperature under each type of tipping point possibility using expected draws. We simulate a path that happens to never cross a threshold in order to see how the modeled policymaker adjusts to the possibility over time.

raises the optimal carbon tax relative to a world without tipping point possibilities. In contrast, the carbon sink tipping point imposes a substantial welfare loss by increasing the time for which emissions affect the atmospheric stock of CO_2. It therefore makes future CO_2 levels higher than they otherwise would have been, increasing future damages for a given emission path. However, because the additional welfare loss occurs far enough in the future and because temperature change is concave in CO_2, the change in welfare due to marginal changes in emissions is similar in the pre- and post-threshold worlds. Optimal post-threshold emissions therefore drive CO_2 concentrations to similar levels regardless of the CO_2 level at the time the threshold is crossed. Total welfare loss and the hazard effect are therefore relatively constant over time with the carbon sink tipping points, and the feedback tipping points' hazard effects can grow substantially larger over time. Therefore, while the optimal near-term carbon tax can be greater under the carbon sink tipping points, the cumulative reduction in CO_2 and temperature can be greater under the feedback tipping points because they make the optimal tax increase faster over time.

Recognizing the present inability of climate science to provide a probability distribution for the temperature threshold, we consider the implications of more and less diffuse priors for the threshold's location and of aversion to the ambiguity in the threshold's distribution. Figure 5 plots the year 2015 optimal carbon tax, the peak CO_2 concentration, and the peak temperature for values of $\bar{T}$ between 3°C and 9°C, with all calculations still being for optimal policy paths conditional on not having crossed the threshold. As the upper bound $\bar{T}$ increases, optimal policy converges asymptotically to the no-threshold policy. A more diffuse prior on the threshold's location reduces the importance of the second and third lines in equation (3) by reducing both the hazard rate and its derivative. However, lowering $\bar{T}$ from its base case value has a stronger effect on optimal policy than does raising $\bar{T}$. For low values of $\bar{T}$, the hazard rate is a steeper function of emissions and realized temperatures can approach regions with a high hazard rate. When $\bar{T} = 3$°C, the optimal near-term carbon tax rises as high as \$XX/tCO_2, with temperature peaking at only XX°C. The policymaker anticipates that a future high hazard effect will justify a high carbon tax (via the third line in equation (3) for future times), which increases the benefits of reducing near-term emissions in order to smooth abatement cost over time (via the first line in equation (3) for the current time).

Finally, we turn to the effect of ambiguity aversion on optimal policy. Figure 6 varies γ from slightly ambiguity-loving ($\gamma = 0 < 2 = \eta$) to extremely ambiguity-averse ($\gamma = 100 > 2 = \eta$). We find that optimal policy is not highly sensitive to the policymaker's level of ambiguity aversion. The near-term carbon tax varies by less than \$2/tCO_2 across the modeled range of γ, the peak CO_2 concentration varies by less than 20 ppm, and peak temperature varies by less than 0.2°C. As noted earlier, the small annual hazard rate means that ambiguity aversion affects the optimal carbon tax primarily by altering the hazard effect.[11] Equation (5) showed that ambiguity aversion changes the hazard effect approximately in proportion to $h - 0.5$, to the total welfare loss from crossing the threshold, and to the measure of absolute ambiguity aversion (near 3 for $\gamma = 100$). The smallness of the annual hazard rate that minimizes pessimism bias therefore strengthens the influence of ambiguity aversion on the hazard effect. Further, ambiguity aversion increases the hazard effect by more for stronger tipping points that have more severe welfare implications. For

[11] Pessimism bias increases the near-term cost of carbon by only XX%.

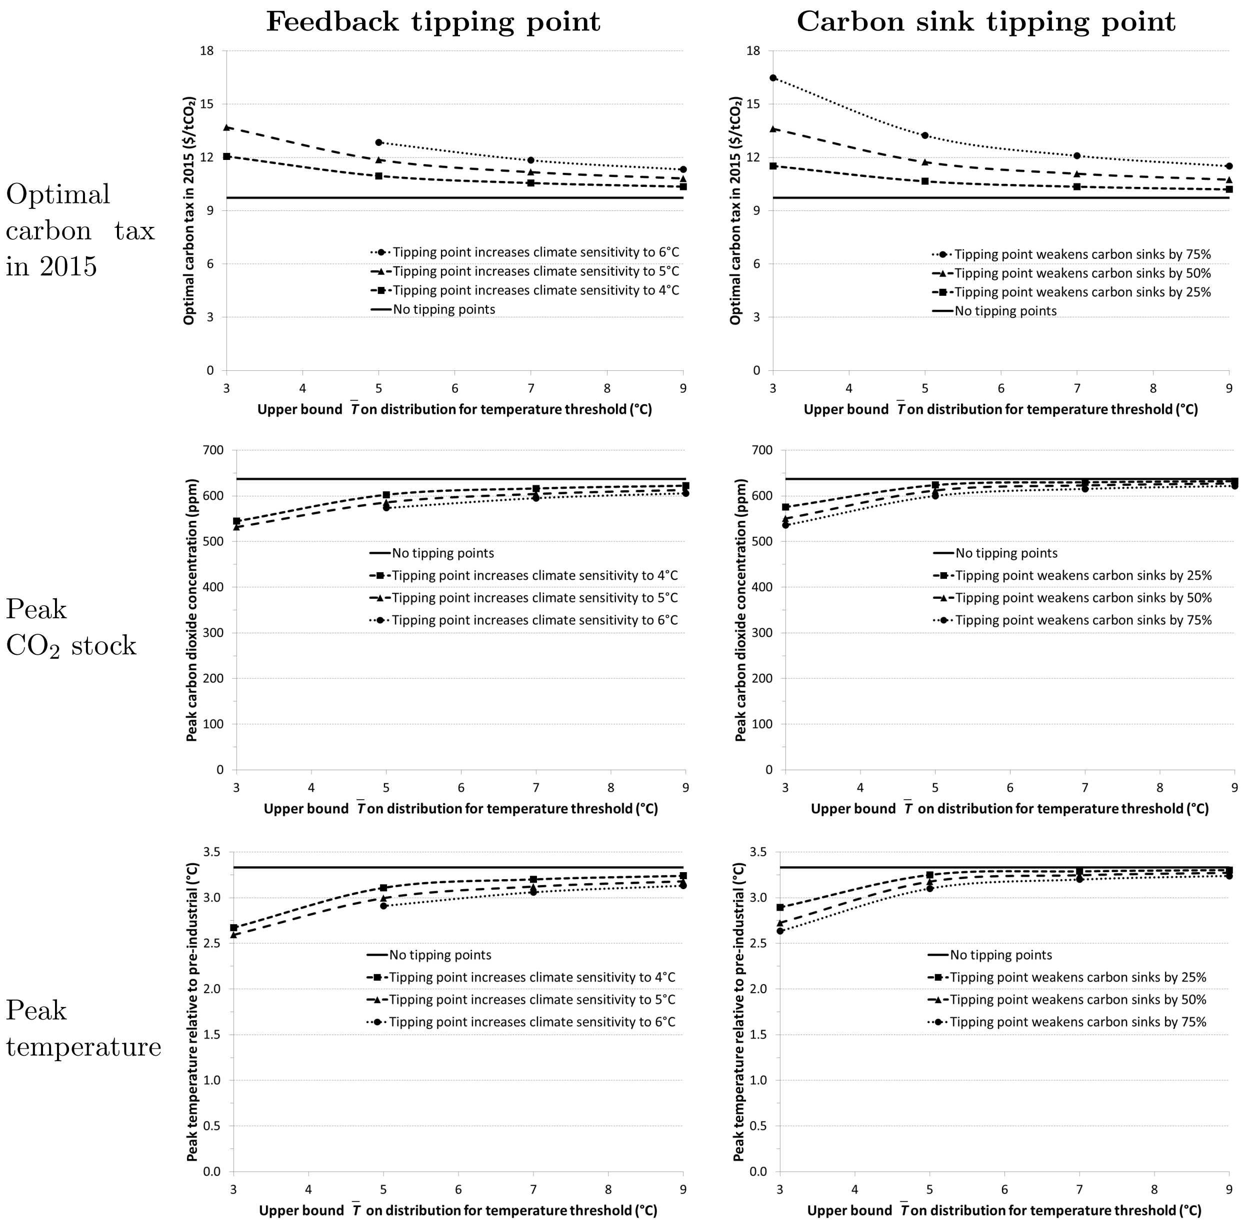

Figure 5: The optimal carbon tax in 2015, the peak CO_2 concentration, and the peak temperature reached for each upper bound $\bar{T}$ for the temperature threshold's distribution. The plotted simulations assume expected draws of the temperature shock and also assume that the tipping point never occurs.

the weakest of the feedback tipping points, the extreme form of ambiguity aversion with $\gamma = 100$ increases the hazard effect by 1% in 2015 and by 15% in 2050, and for the strongest of the feedback tipping points, the extreme form of ambiguity aversion increases the hazard effect by XX% in 2015 and by XX% in 2050. However, the total effect of ambiguity aversion on near-term policy is still small in comparison to the effect of the tipping points themselves. Generalizing the welfare evaluation to acknowledge that tipping point probabilities are, by necessity, ad hoc estimates rather than objective probabilities does not greatly change the importance of tipping points for near-term policy.

6 Conclusions

We have shown how to model economic decisions in the face of irreversible tipping points triggered by policy-dependent thresholds. We analytically decomposed the marginal benefit of a change in a control to show that the dominant channel through which tipping points affect policy will often be the hazard effect. The magnitude of the hazard effect is determined by the degree to which a marginal change in policy affects the probability of crossing a threshold and by the total change in welfare from crossing a threshold. Recognizing that threshold probabilities are often ambiguous, we showed that ambiguity aversion affects the optimal control through pessimism bias (leading the decision-maker to treat the lower-welfare threshold outcome as more likely than it is) and through the hazard effect. Interestingly, ambiguity aversion can actually decrease the hazard effect when the lower-welfare threshold outcome is particularly likely, which means that ambiguity aversion can plausibly lead optimal policy to make the lower-welfare outcome still more likely.

We then integrated an irreversible change of system dynamics into a recursive climate-economy model. We extended the benchmark DICE model to include the endogenous possibility of climatic tipping points, endogenous learning about the temperature threshold that triggers tipping points, endogenous welfare effects from tipping points' occurrence, and a generalized welfare evaluation that allows the policymaker to display ambiguity aversion. We disentangled three components of the optimal carbon tax (social cost of carbon) when tipping points are possible: the effect of emissions on pre-threshold value, the effect of emissions on post-threshold value, and the effect of emissions on the hazard rate. Typical models with smooth dynamics only include the first of these effects. Because our model has only a small probability of crossing a threshold in a given year, it is the hazard effect that almost exclusively determines the effect of tipping point awareness on the optimal carbon tax. Quantitatively, tipping point possibilities can increase the near-term optimal carbon tax by 50%. The precise effect is sensitive to the type of tipping point, to the strength of the tipping point, and to the distribution for the threshold that triggers the tipping point. Carbon sink tipping points tend to have a stronger effect on the near-term carbon tax, but feedback tipping points more strongly increase cumulative abatement by making the carbon tax rise more quickly over time. This result demonstrates the value of explicitly modeling tipping points' effects on system dynamics. Ambiguity aversion raises the optimal carbon tax and does so by an increasing amount over time, but ambiguity attitude is less important for optimal policy than is the type of tipping point considered.

Our numerical conclusions have implications for economic modeling, climate science, and climate

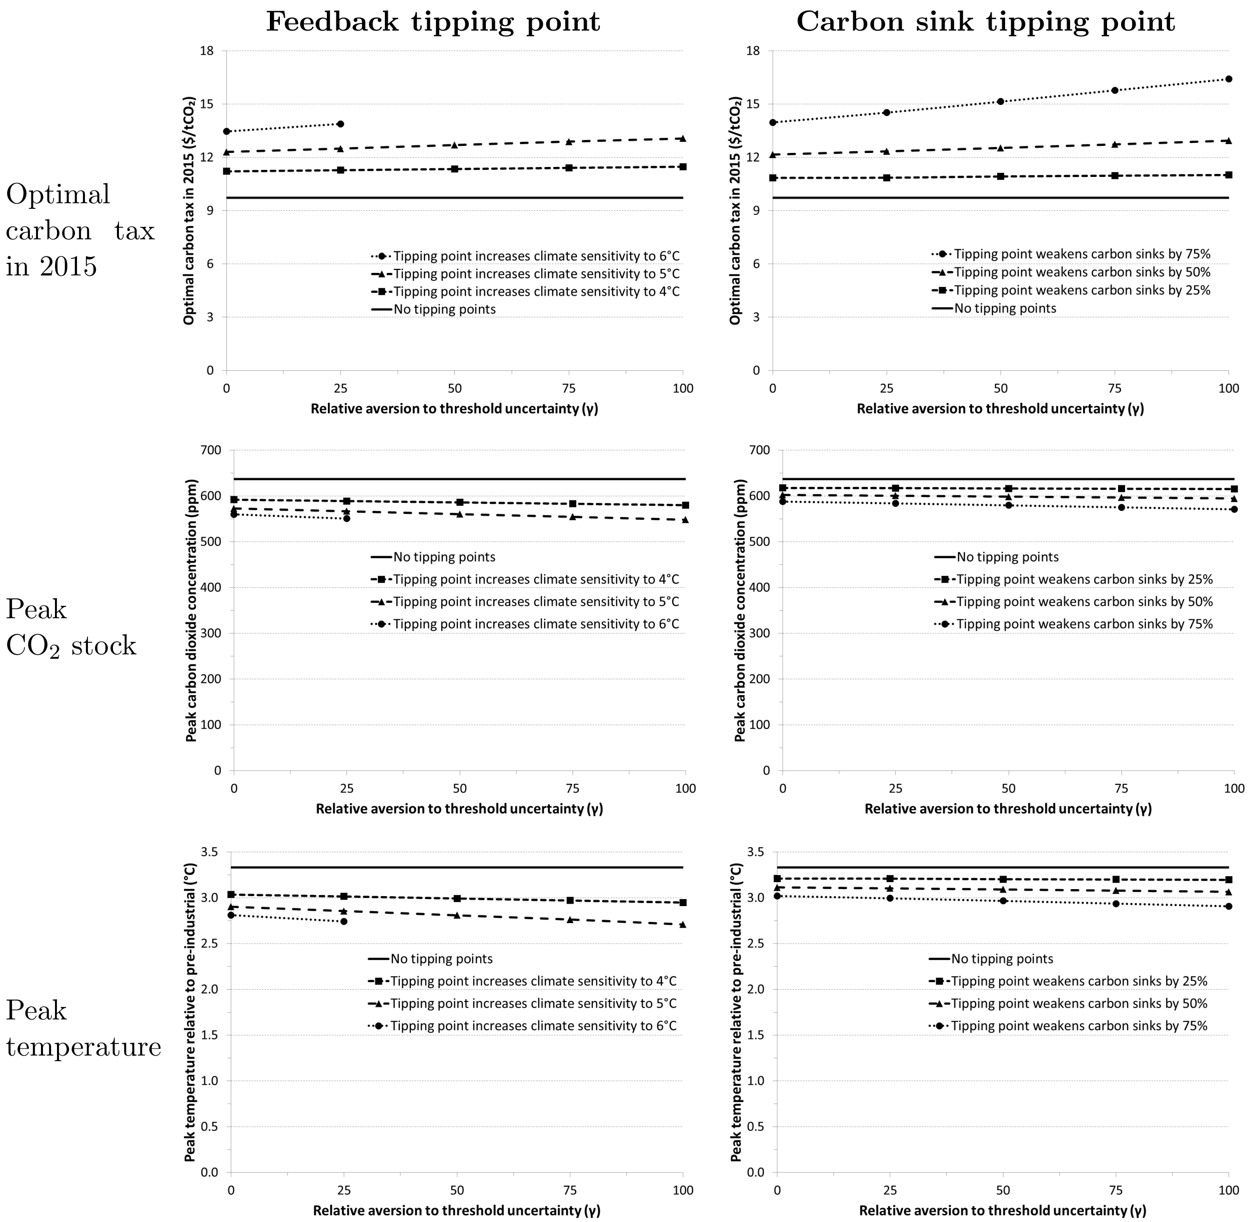

Figure 6: The optimal carbon tax in 2015, the peak CO_2 concentration, and the peak temperature reached for different degrees of aversion to threshold uncertainty (γ). An ambiguity-averse policymaker has $\gamma > 2$. The plotted simulations assume expected draws of the temperature shock and also assume that the tipping point never occurs.

policy. First, economic models of climate change typically assume smooth changes in the climate system. More broadly, nearly all economic models dealing with growth and long-run dynamics assume smoothly changing systems. When models do allow for discontinuous changes, they usually do so by introducing the possibility of a utility penalty or making the change occur exogenously. By showing how feedback and carbon sink tipping points affect the optimal carbon tax in qualitatively different ways, we have demonstrated the value of explicitly modeling the shifts in dynamics that tipping points might provoke. Further, we have also shown that the main effect of tipping points on optimal policy is often due to their endogeneity. It is important to model a tipping point's structural effects rather than reducing it to a predetermined shock to utility, and it is also important to represent the effect of policy decisions on the chance that a tipping point occurs.

Second, our work is a call to the climate sciences to improve knowledge about both the effects of tipping points on system dynamics and the types of temperature paths that trigger them. We have shown that different anticipated changes in dynamics can have very different effects on the optimal carbon tax. We place an economic value on scientific information related to tipping points and open the door to more comprehensive integrated assessments of abrupt climate transitions. However, these more comprehensive assessments depend on the ability of climate science to constrain the probability of crossing thresholds at temperatures close to those of today and to translate the effects of tipping points into the reduced climate models used in climate-economy integrated assessment models.

Third, our findings back up the widespread supposition that the existence of tipping points in the climate system should have a strong influence on our current policy decisions. Moreover, we provide a quantitative basis for adjusting policy for the possibility of tipping points. Integrated assessment models are the main quantitative input into regulations putting a price on carbon. Past compilations of models' estimates described sets of models that jointly omitted climatic features we have shown to be highly relevant. Much work remains to make models' representation of tipping points more realistic, but we have demonstrated how to fully endogenize tipping point possibilities and have provided a first assessment of their effect on policy. Further, by demonstrating the importance of these features absent from previous models, our work also suggests that truly optimal climate policy considers not just the harm from additional emissions in a particular model but also the prospect that the available models do not adequately capture future changes in the climate and the economy.

7 Appendix: The climate-economy model

This appendix provides the complete equations for the climate-economy model extending DICE-2007 from Nordhaus (2008). Table 1 provides the numerical parameterization. The pre-threshold value function is:[12]

$$V_{\psi=0}(k_t,M_t,T_t,t) = \max_{c_t,\mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \frac{\beta_t}{1-\eta} \int \Bigg[\left[1-h(T_t,T_{t+1})\right] \left[(1-\eta)V_{\psi=0}(k_{t+1},M_{t+1},T_{t+1},t+1)\right]^{\frac{1-\gamma}{1-\eta}} + h(T_t,T_{t+1}) \left[(1-\eta)V_{\psi=1}(k_{t+1},M_{t+1},T_{t+1},t+1)\right]^{\frac{1-\gamma}{1-\eta}} \Bigg]^{\frac{1-\eta}{1-\gamma}} d\mathbb{P}$$

subject to

$$k_{t+1} = e^{-(g_{L,t}+g_{A,t})} \left[(1-\delta_k)k_t + (1-\Psi_t \mu_t^{a_2}) \frac{Y_{gross}}{1+D_t} - c_t \right] \quad \text{(Effective capital)}$$

$$M_{t+1} = e_t + M_t \left[b_{11} + b_{21} \left(b_{12} + (b_{22} + b_{32}b_{23})\alpha_B + b_{32}b_{33}\alpha_O \right) \right] \quad (\mathrm{CO_2})$$

$$T_{t+1} = T_t + C_T \left[F(M_{t+1},t+1) - \frac{f}{s} T_t - (1-\alpha_T)\, C_O\, T_t \right] \quad \text{(Temperature)}$$

$$c_t + \Psi_t \mu_t^{a_2} \frac{Y_{gross}}{1+D_t} \leq \frac{Y_{gross}}{1+D_t} \quad \text{(Output constraint)}$$

$$\mu_t \leq 1 \quad \text{(Non-negativity constraint for emissions)}$$

with the state variables being effective capital k_t, atmospheric CO_2 M_t, cumulative temperature change T_t, and time t.[13] The controls are consumption c_t, abatement μ_t, and, as a residual, investment. Welfare in a given period is the sum of immediate utility $u(c_t) = c_t^{1-\eta}/(1-\eta)$ and the discounted expectation of future welfare. The parameter η is the Arrow-Pratt measure of relative risk aversion, and η^{-1} gives the intertemporal elasticity of substitution. The parameter γ measures aversion to tipping point uncertainty, as described in Section 4. The constraints prevent the decision-maker from using more than the output available after accounting for damages and from abating more than 100% of emissions in a period. When the output constraint is slack, we have positive capital investment, and when the abatement constraint is slack, economic activity produces some CO_2 emissions that are not abated.

Capital depreciates at rate δ_k, and capital investment comes from any available output not allocated to the control variables of consumption c_t and abatement μ_t. The exogenous variable Ψ_t and parameter a_2 determine the cost of abating the chosen fraction μ_t of emissions. The term outside the brackets in the capital transition equation adjusts for the growth of labor and technology to keep capital in effective terms. Gross output Y_{gross} is a function of the capital stock:

$$Y_{gross} = k_t^{\kappa} \ . \quad \text{(Gross output)}$$

The parameter κ gives the capital elasticity in a Cobb-Douglas production function. Climate

[12]Normalizing the discount factor, the value function, and some of the equations of motion makes the objective equivalent to maximizing population-weighted per capita consumption $L_t u\,(C_t/L_t)$ as in DICE. The discount factor β_t captures a rate of pure time preference of 1.5% as in DICE-2007. It also adjusts for population growth and technological progress, a step that is part of the normalization yielding a measurement of consumption in effective labor units.

[13]We solve the model using a transformation mapping the infinite time horizon to the unit interval.

damages D_t reduce gross output in accord with the total temperature change:

$$D_t = b_2 \left(\epsilon_t T_t\right)^{b_3} \ , \tag{Damages}$$

where the independent, normally distributed multiplicative shock ϵ_t has probability measure $\mathbb{P}$.[14] Optimal policy adjusts current controls in anticipation of possible future shocks, and a given period's realized shock affects the residual output allocated to investment. We calibrate the mean-1 shock to the years 1881-2010 in the NASA Goddard Institute for Space Studies (GISS) Surface Temperature Analysis dataset.[15] We take expected temperature in each year to be the mean of the surrounding 10 years' realized temperatures. The realized standard deviation of the resulting time series of multiplicative shocks is 0.0068.[16] This multiplicative noise captures period-to-period temperature variability that makes extreme outcomes more likely as CO_2 increases.

The carbon dynamics in DICE-2007 are determined by a transition matrix governing the flow between the atmospheric stock (stock 1), the combined biosphere and shallow ocean stock (stock 2), and the deep ocean stock (stock 3). We approximate the combined biosphere and shallow ocean stock as a fraction α_B of the atmospheric stock and the deep ocean stock as a fraction α_O of the atmospheric stock, where α_B and α_O are interpolated functions of time and CO_2 based on output from several runs of DICE-2007. The parameter b_{12} gives the fraction of atmospheric CO_2 absorbed by land and ocean sinks over a single timestep and is reduced by carbon sink tipping points as described in Section 4. The parameter $b_{11} = 1 - b_{12}$ determines the fraction of CO_2 that remains in the atmosphere from period to period, with the remaining terms in the CO_2 transition equation together governing the transfer of carbon from land and ocean sinks back into the atmosphere. Time t emissions e_t are given by:

$$e_t = \sigma_t (1 - \mu_t) Y_{gross} + B_t \ . \tag{Emissions}$$

The exogenous variable σ_t is the emission intensity of gross output and B_t gives exogenous CO_2 emissions from non-industrial sources such as land use change.

DICE-2007 determines time t surface temperature from the stock of CO_2, from surface temperature in the previous period, and from the previous period's difference between the surface temperature and the deep ocean temperature. We represent the deep ocean temperature as a fraction α_T of surface temperature, where α_T is an interpolated function of time and temperature based on output from several runs of DICE-2007. Forcing $F(M_t, t)$ measures the additional energy (W m^{-2}) trapped at the earth's surface by greenhouse gases and other atmospheric agents. Forcing is concave in CO_2:

$$F(M_t, t) = f \ \ln\left(M_t / M_{pre}\right) + EF_t \ , \tag{Forcing}$$

where f is the forcing from doubled CO_2 and EF_t gives the time t exogenous (non-CO_2) forcing. The parameter s in the temperature transition equation is climate sensitivity, or the equilibrium temperature change from doubling CO_2 concentrations. This parameter is altered by feedback

[14] In the context of Section 2, the shock ϵ_t is a diagonal matrix with the only diagonal element not equal to 1 being in the row corresponding to the temperature state variable.

[15] Available at http://data.giss.nasa.gov/gistemp/.

[16] We implement the continuous distribution numerically using a Gauss-Legendre quadrature rule with 8 nodes.

Table 1: Parameterization of the transition equations. Several values are rounded, and C_T and δ_κ vary slightly over time in order to reproduce the results from DICE-2007 with an annual timestep.

Parameter	Value	Description
A_0	0.027	Initial production technology
$g_{A,0}$	0.009	Initial annual growth rate of production technology
δ_A	0.001	Annual rate of decline in growth rate of production technology
L_0	6514	Population in 2005 (millions)
L_∞	8600	Asymptotic population (millions)
δ_L	0.035	Annual rate of convergence of population to asymptotic value
σ_0	0.13	Initial emission intensity before emission reductions (GtC/output)
$g_{\sigma,0}$	-0.0073	Initial annual growth rate of emission intensity
δ_σ	0.003	Annual change in growth rate of emission intensity
a_0	1.17	Cost of backstop technology in 2005 ($1000/tC)
a_1	2	Ratio of initial backstop cost to final backstop cost
a_2	2.8	Abatement cost exponent
g_Ψ	-0.005	Annual growth rate of backstop cost
B_0	1.1	Initial non-industrial CO_2 emissions (GtC/y)
g_B	-0.01	Annual growth rate of non-industrial emissions
EF_0	-0.06	Initial forcing from non-CO_2 greenhouse gases (W m^{-2})
EF_{100}	0.30	Year 2105 forcing from non-CO_2 greenhouse gases (W m^{-2})
κ	0.3	Capital elasticity in Cobb-Douglas production function
δ_κ	0.06	Annual depreciation rate of capital
b_2	0.0028	Coefficient on temperature in the damage function
b_3	2	Exponent on temperature in the damage function
s	3	Climate sensitivity (°C)
f	3.8	Forcing from doubled CO_2 (W m^{-2})
M_{pre}	596.4	Pre-industrial atmospheric CO_2 (GtC)
C_T	0.03	Translation of forcing into temperature change
C_O	0.3	Translation of surface-ocean temperature gradient into forcing
b_{11},b_{12}	0.978,0.023	Transfer coefficients for carbon from the atmosphere
b_{21},b_{22}, b_{23}	0.011,0.983,0.005	Transfer coefficients for carbon from the combined biosphere and shallow ocean stock
b_{32},b_{33}	0.0003,0.9997	Transfer coefficients for carbon from the deep ocean
ρ	0.015	Annual rate of pure time preference
η	2	Relative risk aversion (also aversion to intertemporal substitution)
k_0	$137/(A_0 L_0)$	Initial effective capital, with initial capital stock of 137 US$trillion
M_0	808.9	Initial atmospheric CO_2 (GtC)
T_0	0.73	Initial surface temperature (°C)

tipping points as described in Section 4. Finally, the parameter C_O determines how a temperature gradient between the surface and the deep ocean affects forcing at the surface, and the parameter C_T controls the speed with which aggregate forcing changes temperature.

We implement our model with an annual timestep, while DICE-2007 uses a decadal timestep. We therefore adjust all transition equations from those in DICE by calculating the parameter values that would reproduce the state variables' paths if the transition equations were instead applied to an annual timestep with constant policies over the decade. Our model replicates DICE's results when run with DICE's policy path or when optimizing with a 10-year timestep as in DICE. However, when optimizing with an annual timestep and without tipping point possibilities, the policymaker's ability to smooth emissions within a decade leads to our peak CO_2 level being about 30 ppm lower than in DICE, our maximum temperature being about 0.15°C lower, and our year 2015 social cost of carbon being about \$4/tCO_2 lower.

The transition equations for the exogenous variables are as follows. In each case, $t = 0$ corresponds to the year 2005.

$$A_t = A_0 \exp\left[\frac{g_{A,0}}{\delta_A}\left(1 - e^{-t\delta_A}\right)\right] \qquad \text{(Production technology)}$$

$$g_{A,t} = g_{A,0} e^{-t\delta_A} \qquad \text{(Growth rate of production technology)}$$

$$L_t = L_0 + (L_\infty - L_0)\left(1 - e^{-t\delta_L}\right) \qquad \text{(Labor)}$$

$$g_{L,t} = \delta_L \left[\frac{L_\infty}{L_\infty - L_0} e^{t\delta_L} - 1\right]^{-1} \qquad \text{(Growth rate of labor)}$$

$$\beta_t = \exp\left(-\rho + (1-\eta) g_{A,t} + g_{L,t}\right) \qquad \text{(Effective discount factor)}$$

$$\sigma_t = \sigma_0 \exp\left[\frac{g_{\sigma,0}}{\delta_\sigma}\left(1 - e^{-t\delta_\sigma}\right)\right] \qquad \text{(Uncontrolled emissions per output)}$$

$$\Psi_t = \frac{a_0 \sigma_t}{a_2}\left(1 - \frac{1 - e^{t g_\Psi}}{a_1}\right) \qquad \text{(Abatement cost factor)}$$

$$B_t = B_0 e^{t g_B} \qquad \text{(Non-industrial } CO_2 \text{ emissions)}$$

$$EF_t = EF_0 + 0.01(EF_{100} - EF_0)\min\{t, 100\} \qquad \text{(Non-} CO_2 \text{ forcing)}$$

The computational challenge in solving the model lies not in finding the optimal actions for a given value function but in determining the value functions that satisfy the relations in equations (1) and (2) (see Kelly and Kolstad, 1999, 2001). We begin with a guess for the value function and a set of Chebychev nodes in the four-dimensional state space. We then use the initial guess for the continuation value to find each node's optimal controls c_t^* and μ_t^* and optimal value. Knowing the optimal value at each Chebychev node, we approximate the value function across the rest of the state space using a set of Chebychev basis polynomials. We repeat the process using this approximated value function as the new initial guess, with iteration continuing until the coefficients of the value approximant's basis functions change by less than 0.0001. The pre-threshold value function is smooth over the relevant state space because the hazard rate is continuous in the state variables.

References

Alley, R. B., J. Marotzke, W. D. Nordhaus, J. T. Overpeck, D. M. Peteet, R. A. Pielke, R. T. Pierrehumbert, P. B. Rhines, T. F. Stocker, L. D. Talley, and J. M. Wallace (2003). "Abrupt climate change." Science 299(5615): 2005–2010. doi:10.1126/science.1081056.

Archer, D. (2007). "Methane hydrate stability and anthropogenic climate change." Biogeosciences 4(4): 521–544.

Azariadis, C. and A. Drazen (1990). "Threshold externalities in economic development." The Quarterly Journal of Economics 105(2): 501–526. doi:10.2307/2937797.

Brock, W. and D. Starrett (2003). "Managing systems with non-convex positive feedback." Environmental and Resource Economics 26(4): 575–602. doi:10.1023/B:EARE.0000007350.11040.e3.

Brozović, N. and W. Schlenker (2011). "Optimal management of an ecosystem with an unknown threshold." Ecological Economics 70(4): 627–640. doi:10.1016/j.ecolecon.2010.10.001.

Camerer, C. and M. Weber (1992). "Recent developments in modeling preferences: Uncertainty and ambiguity." Journal of Risk and Uncertainty 5(4): 325–370. doi:10.1007/BF00122575.

Crost, B. and C. P. Traeger (2010). "Risk and aversion in the integrated assessment of climate change." CUDARE Working Paper 1104, University of California, Berkeley.

Davig, T. and E. M. Leeper (2007). "Generalizing the Taylor principle." American Economic Review 97(3): 607–635. doi:10.1257/aer.97.3.607.

de Zeeuw, A. and A. Zemel (2011). "Regime shifts and uncertainty in pollution control." Working paper.

Eglin, T., P. Ciais, S. L. Piao, P. Barre, V. Bellassen, P. Cadule, C. Chenu, T. Gasser, C. Koven, M. Reichstein, and P. Smith (2010). "Historical and future perspectives of global soil carbon response to climate and land-use changes." Tellus B 62(5): 700–718. doi:10.1111/j.1600-0889.2010.00499.x.

Ellison, G. and D. Fudenberg (2003). "Knife-edge or plateau: When do market models tip?" Quarterly Journal of Economics 2003(4): 1249–78. doi:10.1162/003355303322552793.

Ellsberg, D. (1961). "Risk, ambiguity, and the Savage axioms." The Quarterly Journal of Economics 75(4): 643–669.

Gjerde, J., S. Grepperud, and S. Kverndokk (1999). "Optimal climate policy under the possibility of a catastrophe." Resource and Energy Economics 21(3-4): 289–317. doi:10.1016/S0928-7655(99)00006-8.

Greenstone, M., E. Kopits, and A. 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 No. 16913.

Guo, X., J. Miao, and E. Morellec (2005). "Irreversible investment with regime shifts." Journal of Economic Theory 122(1): 37–59. doi:10.1016/j.jet.2004.04.005.

Hall, D. C. and R. J. Behl (2006). "Integrating economic analysis and the science of climate instability." Ecological Economics 57(3): 442–465. doi:10.1016/j.ecolecon.2005.05.001.

Hansen, J., M. Sato, P. Kharecha, D. Beerling, R. Berner, V. Masson-Delmotte, M. Pagani, M. Raymo, D. L. Royer, and J. C. Zachos (2008). "Target atmospheric CO2: Where should humanity aim?" The Open Atmospheric Science Journal 2: 217–231. doi:10.2174/1874282300802010217.

Heal, G. (1984). "Interactions between economy and climate: A framework for policy design under uncertainty." In V. K. Smith and A. D. White, eds., "Advances in Applied Microeconomics," volume 3, pp. 151–168. Greenwich, CT: JAI Press.

Huntingford, C., R. A. Fisher, L. Mercado, B. B. Booth, S. Sitch, P. P. Harris, P. M. Cox, C. D. Jones, R. A. Betts, Y. Malhi, G. R. Harris, M. Collins, and P. Moorcroft (2008). "Towards quantifying uncertainty in predictions of Amazon 'dieback'." Philosophical Transactions of the Royal Society B: Biological Sciences 363(1498): 1857 –1864. doi:10.1098/rstb.2007.0028.

Katz, M. L. and C. Shapiro (1994). "Systems competition and network effects." The Journal of Economic Perspectives 8(2): 93–115.

Keller, K., B. M. Bolker, and D. F. Bradford (2004). "Uncertain climate thresholds and optimal economic growth." Journal of Environmental Economics and Management 48(1): 723–741. doi:10.1016/j.jeem.2003.10.003.

Kelly, D. L. and C. D. Kolstad (1999). "Bayesian learning, growth, and pollution." Journal of Economic Dynamics and Control 23(4): 491–518. doi:10.1016/S0165-1889(98)00034-7.

Kelly, D. L. and C. D. Kolstad (2001). "Solving infinite horizon growth models with an environmental sector." Computational Economics 18(2): 217–231. doi:10.1023/A:1021018417052.

Keynes, J. M. (1921). A Treatise on Probability. London: Macmillan and Co.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). "A smooth model of decision making under ambiguity." Econometrica 73(6): 1849–1892.

Klibanoff, P., M. Marinacci, and S. Mukerji (2009). "Recursive smooth ambiguity preferences." Journal of Economic Theory 144(3): 930–976. doi:10.1016/j.jet.2008.10.007.

Kriegler, E., J. W. Hall, H. Held, R. Dawson, and H. J. Schellnhuber (2009). "Imprecise probability assessment of tipping points in the climate system." Proceedings of the National Academy of Sciences 106(13): 5041–5046. doi:10.1073/pnas.0809117106.

Lange, A. and N. Treich (2008). "Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions." Climatic Change 89(1): 7–21. doi:10.1007/s10584-008-9401-5.

Le Quéré, C., C. Rodenbeck, E. T. Buitenhuis, T. J. Conway, R. Langenfelds, A. Gomez, C. Labuschagne, M. Ramonet, T. Nakazawa, N. Metzl, N. Gillett, and M. Heimann (2007). "Saturation of the Southern Ocean CO2 sink due to recent climate change." Science 316(5832): 1735–1738. doi:10.1126/science.1136188.

Leach, A. J. (2007). "The climate change learning curve." Journal of Economic Dynamics and Control 31(5): 1728–1752. doi:10.1016/j.jedc.2006.06.001.

Lenton, T. M., H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber (2008). "Tipping elements in the Earth's climate system." Proceedings of the National Academy of Sciences 105(6): 1786–1793. doi:10.1073/pnas.0705414105.

Mäler, K.-G., A. Xepapadeas, and A. de Zeeuw (2003). "The economics of shallow lakes." Environmental and Resource Economics 26(4): 603–624. doi:10.1023/B:EARE.0000007351.99227.42.

Millner, A., S. Dietz, and G. Heal (2010). "Ambiguity and climate policy." National Bureau of Economic Research Working Paper Series No. 16050.

Nævdal, E. (2006). "Dynamic optimisation in the presence of threshold effects when the location of the threshold is uncertain—with an application to a possible disintegration of the Western Antarctic Ice Sheet." Journal of Economic Dynamics and Control 30(7): 1131–1158. doi:10.1016/j.jedc.2005.04.004.

Nævdal, E. and M. Oppenheimer (2007). "The economics of the thermohaline circulation—A problem with multiple thresholds of unknown locations." Resource and Energy Economics 29(4): 262–283. doi:10.1016/j.reseneeco.2007.01.003.

Nordhaus, W. D. (1992). "An optimal transition path for controlling greenhouse gases." Science 258(5086): 1315–1319.

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

Overpeck, J. T. and J. E. Cole (2006). "Abrupt change in Earth's climate system." Annual Review of Environment and Resources 31(1): 1–31. doi:10.1146/annurev.energy.30.050504.144308.

Polasky, S., A. de Zeeuw, and F. Wagener (2011). "Optimal management with potential regime shifts." Journal of Environmental Economics and Management 62(2): 229–240. doi:10.1016/j.jeem.2010.09.004.

Ramanathan, V. and Y. Feng (2008). "On avoiding dangerous anthropogenic interference with the climate system: Formidable challenges ahead." Proceedings of the National Academy of Sciences 105(38): 14245 –14250. doi:10.1073/pnas.0803838105.

Schaefer, K., T. Zhang, L. Bruhwiler, and A. P. Barrett (2011). "Amount and timing of permafrost carbon release in response to climate warming." Tellus B 63(2): 165–180. doi:10.1111/j.1600-0889.2011.00527.x.

Schelling, T. (1971). "Dynamic models of segregation." The Journal of Mathematical Sociology 1(2): 143–186. doi:10.1080/0022250X.1971.9989794.

Smith, J. B., S. H. Schneider, M. Oppenheimer, G. W. Yohe, W. Hare, M. D. Mastrandrea, A. Patwardhan, I. Burton, J. Corfee-Morlot, C. H. D. Magadza, H. Fussel, A. B. Pittock, A. Rahman, A. Suarez, and J. van Ypersele (2009). "Assessing dangerous climate change through an update of the Intergovernmental Panel on Climate Change (IPCC) "reasons for concern"." Proceedings of the National Academy of Sciences 106(11): 4133–4137. doi:10.1073/pnas.0812355106.

Traeger, C. P. (2010). "Subjective risk, confidence, and ambiguity." CUDARE Working Paper 1103, University of California, Berkeley.

Tsur, Y. and A. Zemel (1996). "Accounting for global warming risks: Resource management under event uncertainty." Journal of Economic Dynamics and Control 20(6-7): 1289–1305. doi:10.1016/0165-1889(95)00900-0.

Wagener, F. O. O. (2003). "Skiba points and heteroclinic bifurcations, with applications to the shallow lake system." Journal of Economic Dynamics and Control 27(9): 1533–1561. doi:16/S0165-1889(02)00070-2.