Ambiguous Tipping Points*

Derek Lemoine[†] & Christian P. 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
& Department of Economics, Postbox 1095 Blindern, 0317 Oslo, Norway,
traeger@berkeley.edu

February 2016

We analyze the policy implications of aversion to Knightian uncertainty (ambiguity) about the possibility of tipping points. We demonstrate two channels through which uncertainty aversion affects optimal policy in the general setting. The first channel relates to the policy's effect on the probability of tipping, and the second channel to its differential impact in the pre- and post-tipping regimes. We then extend a recursive dynamic model of climate policy and tipping points to include uncertainty aversion. Numerically, aversion to Knightian uncertainty in the face of an ambiguous tipping point increases the optimal tax on carbon dioxide emissions, but only by a small amount.

JEL Codes: Q54, D90, D81, C61, H23

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

*We are grateful for comments from Larry Karp, Florian Diekert, Ujjayant Chakravorty, and participants at a number of seminars, including the 2011 NBER Environmental and Energy Economics Summer Institute and the 2015 workshop on “Thresholds, Tipping Points and Random Events in Dynamic Economic Systems” at the University of Tennessee. Some of this material previously circulated in working paper form as “Tipping Points and Ambiguity in the Economics of Climate Change.” Christian Traeger gratefully acknowledges support by the National Science Foundation through the Network for Sustainable Climate Risk Management GEO-1240507.

1 Introduction

Tipping points confront policy makers with the possibility of regime shifts. Policymakers usually have a thin or even nonexistent record of past regime shifts from which to gauge the possibility of triggering such a shift in the future. For instance, European countries deciding whether to tolerate the bankruptcy of a member state must consider the unfamiliar chance of tipping into a new regime of high bond yields and further crises. Political elites deciding whether to expropriate resources must consider the chance of triggering a rare mass uprising. And, in our example of primary interest, policymakers deciding how to regulate greenhouse gas emissions must consider the unexperienced chance of irreversibly tipping the planet into a less favorable climate system.

Many agents appear to dislike Knightian uncertainty (Camerer and Weber 1992), and normative models of decision-making under uncertainty allow for such aversion (Traeger 2010, Cerreia-Vioglio et al. 2011, Gilboa and Marinacci 2013). Our policymaker demonstrates an aversion to low-confidence Bayesian priors in the framework of Traeger (2010), which relates closely to the “smooth” ambiguity model of Klibanoff et al. (2009).¹ We analyze how aversion to Knightian uncertainty alters optimal policy in the presence of tipping points with unknown triggers. We numerically solve for the optimal tax on carbon dioxide emissions in the face of deeply uncertain climate tipping points.²

Lemoine and Traeger (2014) show that potential tipping points affect the optimal level of a policy control through two channels. We show how aversion to (Knightian) uncertainty about the threshold’s location changes each of these channels.³ First, the endogenous probability of tipping implies that the present policy influences the chance of tipping. Lemoine and Traeger (2014) call this channel the marginal hazard effect (MHE)

¹We use the terms “Knightian uncertainty”, “deep uncertainty”, “uncertainty”, and “ambiguity” interchangeably to describe a situation where the underlying probabilities are not known. We allow our decision-maker to exhibit higher aversion to situations with unknown probabilities. The smooth ambiguity model of Klibanoff et al. (2005, 2009) relies on subjective ambiguous priors over objective lotteries. By detaching ambiguity aversion from the particular hierarchical order of priors in the recursive smooth ambiguity model, the framework of Traeger (2010) allows us to separate the ambiguous probability of triggering a tipping point from stochastic shocks whose distribution is objectively known.

²Gjerde et al. (1999), Keller et al. (2004), and Lontzek et al. (2015) also numerically analyze the implications of endogenous climate tipping points. Whereas they model tipping points as directly reducing utility or output, we follow Lemoine and Traeger (2014) in modeling tipping points as directly shifting the underlying dynamics of the climate system. The implications for utility and output then depend on how the policymaker responds to the altered climate dynamics. van der Ploeg (2014) follows Lemoine and Traeger (2014) in modeling a tipping point as changing the dynamics of the carbon cycle.

³Our analysis of uncertainty aversion corresponds to an analysis of risk aversion in a setting in which tipping points are the only stochastic element and Epstein-Zin preferences disentangle risk aversion from intertemporal consumption smoothing motives.

and show that it leads the optimal policy to reduce the likelihood of a harmful tipping point. We show that uncertainty aversion further increases this marginal hazard effect if the tipping hazard is small (and decreases the MHE if the probability of tipping is large). A rule of thumb is that uncertainty aversion changes the MHE in a direction that reduces Knightian uncertainty. In our climate change application, the annual probability of tipping is small, so uncertainty aversion increases the contribution of the MHE to the optimal tax on carbon dioxide emissions.

Second, the marginal effect of today's policy on future welfare depends on whether a tipping point happens to occur. Lemoine and Traeger (2014) call this channel the differential welfare impact (DWI): it is proportional to the difference in the control's welfare impact in the pre- and post-tipping regimes. If the marginal welfare benefit of stronger policy is greatest in the post-tipping regime, then uncertainty aversion tends to strengthen policy through this second channel. In our climate example, tipping into a runaway climate can imply that strengthening policy generates higher payoffs in the post-tipping world. This effect would dominate if tipping risk was exogenous. However, the endogeneity of the tipping hazard adds an additional value to emission reductions in the pre-threshold regime, which can outweigh the increased harm from carbon dioxide emissions in the post-threshold world (i.e., in a world with a runaway climate). Our numeric application finds that uncertainty aversion makes the DWI reduce the optimal tax on carbon dioxide emissions. However, the reduction in the DWI is small in comparison to the effect of uncertainty aversion on the marginal hazard effect. Overall, uncertainty aversion increases the optimal tax on carbon dioxide emissions.

Our work sits at the intersection of four literatures. First, a primarily theoretical literature investigates the implications of uncertainty aversion for optimal savings and portfolio allocations. Our DWI contribution relates closely to a self-insurance motive identified in these settings (Gollier 2011, Alary et al. 2013). Our MHE contribution relates closely to a self-protection motive (Snow 2011, Alary et al. 2013, Maccheroni et al. 2013). Traeger (2011) relates these uncertainty aversion effects to the precautionary savings arguments under risk. Our framework differs in two ways from this literature. First, it is more complex because self-insurance and self-protection interact (including over time), and because our discrete regime shift complicates the trade-off for self-insurance. Second, our setting is simpler in that mitigation effort only affects the uncertain tipping hazard, whereas much of this literature has intertwined self-protection against uncertainty with self-protection against standard risk. For the MHE effect, our findings are closest to Snow (2011), who finds that uncertainty aversion always increases self-protection, which holds in

our setting under an approximation for small to moderately large hazard rates. As Alary et al. (2013) explain, this result arises in our setting because mitigation effort reduces uncertainty. The self-insurance motive in our setting, captured by the DWI, differs more significantly from the analysis in this literature because of our discrete regime shift.

Second, the literature has considered the implications of Knightian uncertainty in the context of climate change. In a stylized two-period example, Lange and Treich (2008) show that aversion to Knightian uncertainty about damages from climate change reduces optimal emissions. Millner et al. (2013) analyze the implications of scientific ambiguity about the effect of greenhouse gas emissions on global temperature. They consider how uncertainty aversion affects the welfare under exogenously defined emission policies. They find that the effect is analytically ambiguous but numerically works to increase the value of emission policies that limit warming. Traeger (2014) analyzes the social discount rate of adaptation and mitigation projects under ambiguity about the correlation between economic growth and the projects payoff and finds that ambiguity aversion has a negligible impact. Berger et al. (2015) analyze the implications of ambiguity about the endogenous probability that climate change will reduce consumption. Other papers use Knightian uncertainty to motivate robust control approaches to emission policy (Hennlock 2009, Anderson et al. 2014, Li et al. 2014, Rudik 2015). Our work relates most closely to Jensen and Traeger (2013, 2014), who analytically and numerically analyze the effects of Epstein-Zin preferences and smooth ambiguity aversion on optimal emission policy in a setting without tipping points.

Third, our numerical exploration of uncertainty aversion in the context of tipping points extends a recent literature that uses recursive dynamic models to explore the implications of uncertainty and learning for the optimal tax on carbon dioxide emissions. This literature has analyzed the implications of uncertainty about warming (Kelly and Kolstad 1999, Leach 2007, Jensen and Traeger 2013, Hwang et al. 2014, Kelly and Tan 2015), about economic growth (Jensen and Traeger 2014), about damages from climate change (Crost and Traeger 2013, 2014, Cai et al. 2013, Rudik 2015, Fitzpatrick and Kelly 2015), and about tipping points under standard risk (Lemoine and Traeger 2014).

Finally, the environmental economics literature has a long history of discussing theoretical aspects of pollution-triggered regimes shifts. Clarke and Reed (1994) discuss a regime shift triggered by a stochastic process. Heal (1984) and Tsur and Zemel (1996) introduce epistemological uncertainty where the policymaker does not know the underlying threshold, as in our setting. Tsur and Zemel (2009) and de Zeeuw and Zemel (2012) emphasize that endogenous tipping risk can make the policymaker more cautious, and

Nævdal (2006) analyzes the case where the tipping point’s trigger is a combination of different states. Diekert (2015) discusses how a tipping hazard affects the Nash equilibrium in a non-cooperative resource game, and Barrett and Dannenberg (2012) and Barrett and Dannenberg (2014) discuss how threshold uncertainty challenges cooperative solutions in (stylized) climate threshold experiments.

The next section describes the theoretical setting and the determinants of optimal policy. Section 3 analyzes how uncertainty aversion affects the value of increasing a control. Section 4 describes our recursive dynamic model of climate policy, ultimately based on the benchmark DICE-2007 integrated assessment model of Nordhaus (2008). Section 5 reports our numerical results. Section 6 concludes.

2 Setting

This section explains our representation of tipping points as a regime shift triggered with an endogenous probability. Our policymaker faces an intertemporal trade-off between present benefits and future damages. We show how the presence of regime shifts changes this trade-off.

2.1 The tipping point

We analyze optimal policy in the face of irreversible shifts in system dynamics. These regime shifts are triggered upon crossing a threshold in the state space. The policymaker does not know the precise location of the threshold, has a subjective prior over the possible threshold locations, and learns about its location by observing whether tipping occurs. Our setting embeds the tipping points of Lemoine and Traeger (2014) into a setting with Knightian uncertainty.

The hazard rate is the probability of tipping into a new regime within a given period. This hazard rate is a function of prior beliefs about possible threshold locations and the set of possible threshold locations that the system passes through in a given period. The policymaker controls the tipping probability by controlling movement through the state space. This control is imperfect because, first, inertia in the system prevents full and immediate control of the system states and, second, because the system is subject to stochastic shocks. Our policymaker learns that thresholds that have not triggered a tipping point in the past will be safe in the future, and he updates his Bayesian prior governing threshold locations accordingly.

The policymaker is fully rational and forward-looking, and he incorporates the value of learning into his decisions. He solves an infinite-horizon dynamic optimization problem. We denote the time t vector of system states by S_t . In our climate change example, this vector will include capital, temperature, carbon dioxide, and time (to account for exogenous nonstationary processes). The value function $V_\psi(S_t)$ gives the expected welfare obtained from the optimally controlled system. The parameter ψ indicates the regime: $\psi = 0$ indicates the pre-tipping world, whereas $\psi = 1$ indicates the post-tipping world. In our climate change example, the carbon dioxide and temperature dynamics will differ between the pre- and post-tipping regimes.

The policymaker's utility $u(x_t, S_t)$ in a period depends on the system's state and on the control vector x_t . The optimal post-tipping policy and value function solve the standard dynamic programming problem:

$$\begin{aligned} V_1(S_t) = \max_{x_t} & \left\{ u(x_t, S_t) + \beta_t \int V_1(S_{t+1}) d\mathbb{P} \right\} \\ \text{s.t. } & S_{t+1} = g_1(x_t, \epsilon_t, S_t) \\ & x_t \in \Gamma(S_t) . \end{aligned} \tag{1}$$

The function $g_1(x_t, \epsilon_t, S_t)$ characterizes the next period's state as a function of this period's control vector x_t , the stochastic shock vector ϵ_t , and this period's state vector S_t . The set $\Gamma(S_t)$ constrains the control vector x_t . The probability measure $\mathbb{P}$ characterizes the distribution of the independently and identically distributed stochastic shocks ϵ_t . The Bellman equation (1) is a functional equation for the value function V_1 . Its solution determines the maximal value that the policymaker can obtain in a given state of the world if the regime shift has already occurred.

Prior to crossing a threshold, the policymaker distinguishes between the case where the system dynamics remain the same and the case where the system tips. The hazard rate $h(S_t, S_{t+1})$ gives the probability of tipping.⁴ By assumption, we have no record of events that previously triggered the tipping point, and abrupt regime shifts in system dynamics are hard to forecast accurately. The hazard rate of tipping points therefore exhibits deep uncertainty: the probabilities governing tipping points are not confidently known. Deep uncertainty is often associated with greater uncertainty aversion (Camerer and Weber 1992).

⁴In general, the state space contains informational variables that tell the policymaker which part of the state space has already been explored. In our climate application, the policymaker keeps track of the greatest historic temperature.

We use a normatively-founded rational model of deep uncertainty axiomatized in Traeger (2010), which relates closely to the recursive smooth ambiguity model of Klibanoff et al. (2005, 2009).⁵ As Traeger (2010) explains, the model also relates closely to the Epstein and Zin (1989) and Weil (1989) models that disentangle risk attitude from the policymaker's propensity to smooth consumption over time. The long-run risk literature in asset pricing finds that risk aversion is substantially larger than the willingness to smooth consumption over time (e.g., Vissing-Jørgensen and Attanasio 2003, Bansal and Yaron 2004, Nakamura et al. 2013). In these recursive utility models, a concave function f_{unc} captures how the policymaker is more averse to risk than to known changes in consumption over time. This function f_{unc} should be even more concave when policymakers are averse to deep uncertainty and face ambiguous hazards, as in our setting. The main feature making our model one of deep uncertainty is that we apply the aversion function f_{unc} only to the subjectively uncertain tipping hazard and not to the stochastic shocks ϵ_t . In the absence of the stochastic shocks, our results for uncertainty aversion could be interpreted as the effects of risk aversion when disentangled from willingness to substitute consumption over time.

In the pre-threshold regime, the policymaker solves the dynamic programming problem:

$$\begin{aligned}
 V_0(S_t) = \max_{x_t} & \left\{ u(x_t, S_t) + \beta \int f_{unc}^{-1} \left[[1 - h(S_t, S_{t+1})] f_{unc}[V_0(S_{t+1})] \right. \right. \\
 & \left. \left. + h(S_t, S_{t+1}) f_{unc}[V_1(S_{t+1})] \right] d\mathbb{P} \right\} \quad (2) \\
 \text{s.t. } & S_{t+1} = g_0(x_t, \epsilon_t, S_t) \\
 & x_t \in \Gamma(S_t) .
 \end{aligned}$$

With probability h the system tips and V_1 determines the maximal expected future welfare from period $t + 1$ on. With probability $1 - h$, the system dynamics remain unaltered and V_0 characterizes future welfare. The expression under the integral evaluates the future as the nonlinear weighted mean of the welfare derived in the case that the system remains unaltered and in the case that the system tips. The concave weighting function f_{unc} captures uncertainty aversion, as described above. The integral itself takes expectations over the stochastic shocks. In contrast to the distribution governing the chance of a regime shift, we assume that the distribution governing these temperature shocks is well

⁵The normative axiomatization relies on the von Neumann-Morgenstern axioms, but recognizes that a decision-maker can have different confidence in different distributions.

understood, derived from a long record of observations. These shocks are therefore not deeply uncertain (or ambiguous), and their effects are evaluated outside the Knightian uncertainty (or ambiguity) operator f_{unc} .

2.2 The basic policy trade-off

We now identify the channels through which tipping points affect optimal policy. The subsequent section explains how uncertainty aversion alters each channel. The possible existence of a tipping point introduces two new terms into the marginal welfare impact of changing a control. For ease of exposition, we analyze the case where a single state variable determines the chance of crossing the threshold and the constraints $\Gamma(\cdot)$ do not bind. The right-hand side of equation (2) characterizes welfare for an optimal choice of the controls (with optimality denoted by $*$). We evaluate the marginal welfare impact of varying a generic entry e_t of the control vector in the neighborhood of the optimum. In our climate change application, the temperature state variable determines the hazard, and the welfare impact of varying emissions determines the optimal carbon tax. Suppressing all arguments independent of e_t , the value of the optimal policy program is:

$$u(e_t^*) + \beta_t \int \underbrace{f_{unc}^{-1} \left[[1 - h(S_{t+1}(e_t^*))] f_{unc}[V_0(S_{t+1}(e_t^*))] + h(S_{t+1}(e_t^*)) f_{unc}[V_1(S_{t+1}(e_t^*))] \right]}_{V_{eff}(e_t^*)} d\mathbb{P}.$$

The integrand $V_{eff}(e_t^*)$ expresses the value of future periods' optimal policy program in utility units. It is the uncertainty-averse mean of pre- and post-threshold value. The total value aggregates V_{eff} over standard risk and combines it with the immediate utility from current choices. Varying e_t gives the following trade-off characterizing optimal policies:

$$\begin{aligned} \frac{\partial u(e_t^*)}{\partial e_t} = & -\beta \int \left\{ [1 - h(S_{t+1}(e_t^*))] \frac{f'_{unc}[V_0(S_{t+1}(e_t^*))]}{f'_{unc}[V_{eff}(e_t^*)]} \frac{\partial V_0(S_{t+1}(e_t^*))}{\partial S_{t+1}} \frac{\partial S_{t+1}(e_t^*)}{\partial e_t} \right. \\ & + h(S_{t+1}(e_t^*)) \frac{f'_{unc}[V_1(S_{t+1}(e_t^*))]}{f'_{unc}[V_{eff}(e_t^*)]} \frac{\partial V_1(S_{t+1}(e_t^*))}{\partial S_{t+1}} \frac{\partial S_{t+1}(e_t^*)}{\partial e_t} \\ & \left. - \frac{\partial h(S_{t+1}(e_t^*))}{\partial S_{t+1}} \frac{\partial S_{t+1}(e_t^*)}{\partial e_t} \frac{f_{unc}[V_0(S_{t+1}(e_t^*))] - f_{unc}[V_1(S_{t+1}(e_t^*))]}{f'_{unc}[V_{eff}(e_t^*)]} \right\} d\mathbb{P}, \end{aligned} \quad (3)$$

where primes (') indicate derivatives.⁶ We interpret this equation for the case where an increase in e_t raises current utility but decreases expected future welfare. For instance, additional carbon dioxide emissions increase current utility but decrease future welfare by generating higher carbon stocks and temperatures; additional borrowing increases current consumption but also increases future debt; and a political elite's appropriation of resources provides them with current benefits but increases social discontent and hampers growth. By virtue of Bellman's principle of optimality, equation (3) boils the complex intertemporal decision tree down into a two period trade-off. All future periods face a similar trade-off (conditional on not having tipped). These future trade-offs and policy responses are captured in the (pre- and post-threshold) value functions for period $t + 1$.

The left-hand side of equation (3) characterizes the (immediate) benefits from increasing the policy variable. At an optimum, these benefits must balance the expected future costs. The costs are represented by the right-hand side of equation (3) and are subject to uncertainty. The integrand in the first line represents the impact of policy on time $t + 1$ welfare under the pre-threshold regime (i.e., on the pre-threshold continuation value). This impact is composed of the control's impact on the state vector and the effect of the altered state vector on pre-threshold value V_0 , and it is weighted by the probability $(1 - h)$ of staying in the pre-threshold regime. In a world without tipping points (where the hazard rate h is zero), the first line characterizes the full trade-off between current and future welfare.

Lemoine and Traeger (2014) explain that tipping points extend the standard trade-off in two ways.⁷ First, optimal policy now accounts for its ability to influence the endogenous probability of the regime shift. They call this channel the marginal hazard effect (MHE). It is captured in the third line of equation (3). When triggering a tipping point reduces welfare, this channel moves the optimal control in a direction that reduces the probability of tipping. Second, today's policy affects conditions in the shifted regime in the case that tipping happens to occur. Lemoine and Traeger (2014) call this channel the differential welfare impact (DWI). It is captured by the terms with h in the first two lines of equation (3). Its ambiguous sign depends on how changing the control affects welfare in the pre- and post-threshold worlds. We now show how uncertainty aversion changes these two tipping point contributions to optimal policy.

⁶For a multidimensional state space, $\partial V / \partial S_{t+1}$ and $\partial h / \partial S_{t+1}$ denote gradients and $\partial S_{t+1} / \partial e_t$ denotes the vector of state changes caused by the marginal change in e_t . These derivatives are taken with respect to the pre-threshold dynamics g_0 .

⁷In addition, the potential for future tipping points changes the pre-threshold continuation value and so also affects policy via the first line.

3 Uncertainty aversion and optimal policy in the face of a tipping point

This section studies the trade-off in equation (3) to show how uncertainty aversion affects optimal policy in the face of a possible tipping point. First, we explain how uncertainty aversion affects optimal policy for an endogenous tipping hazard by changing the willingness to invest in reducing the hazard of tipping. Second, we explain how uncertainty aversion affects optimal policy under a (possibly exogenous) tipping hazard by accounting for a policy's impact on post-threshold welfare. We use “strengthening” a policy to describe a change that counteracts an externality that makes tipping points more likely. In our climate change application, strengthening policy means reducing the emission control or increasing the carbon tax.

3.1 The policy implications of an endogenous hazard

When the probability of tipping is endogenous, the optimal policy accounts for the marginal impact of the control e_t on the hazard rate. The third line in equation (3) summarizes this endogenous hazard impact:

$$MHE^{total} = - \underbrace{\frac{\partial h(S_{t+1}(e_t^*))}{\partial S_{t+1}}}_{(i)} \underbrace{\frac{\partial S_{t+1}(e_t^*)}{\partial e_t}}_{(ii)} \underbrace{\frac{f_{unc}[V_0(S_{t+1}(e_t^*))] - f_{unc}[V_1(S_{t+1}(e_t^*))]}{f'_{unc}[V_{eff}(e_t^*)]}}_{(iii)}. \quad (4)$$

Lemoine and Traeger (2014) derive the uncertainty-neutral form:

$$MHE^{neutral} = \frac{\partial h(S_{t+1}(e_t^*))}{\partial S_{t+1}} \frac{\partial S_{t+1}(e_t^*)}{\partial e_t} [V_0 - V_1].$$

They call it the *marginal hazard effect* (MHE). As they explain, the MHE is composed of the response of the hazard rate to a change in the state vector (term i), the response of the state vector to a change in the control (term ii), and the total welfare change from switching regimes (term iii). The uncertainty-neutral marginal hazard effect always strengthens the optimal policy if the externality increases the chance of triggering a harmful tipping point.

We analyze how uncertainty aversion affects this marginal hazard effect via term iii in

equation (4):

$$\frac{f_{unc}[V_0] - f_{unc}[V_1]}{f'_{unc}[V_{eff}]} = [V_0 - V_1] \underbrace{\frac{f_{unc}[V_0] - f_{unc}[V_1]}{[V_0 - V_1] f'_{unc}[V_{eff}]}}_{MHE_{multiplier}^{unc}}. \quad (5)$$

The fraction characterizes the contribution of uncertainty aversion. If the fraction is less than unity, then uncertainty aversion reduces the MHE. It is unity for an uncertainty-neutral decision-maker and if $f'_{unc}[V_{eff}] = \frac{f_{unc}[V_0] - f_{unc}[V_1]}{V_0 - V_1}$. By the mean value theorem, the latter equality has to be satisfied for some V_{eff} and, thus, for some hazard $h^{critical}$.⁸ As the hazard falls, V_{eff} increases because this generalized mean of pre- and post-threshold value places less weight on the smaller post-threshold value function V_1 (assuming a harmful tipping point). Given the concavity of f_{unc} , reducing a hazard from the critical value therefore reduces the term $f'_{unc}[V_{eff}]$ in the denominator of equation (5), which makes the uncertainty multiplier larger than unity. Thus, uncertainty aversion amplifies the marginal hazard effect for hazard rates below the critical value. Analogously, uncertainty aversion reduces the marginal hazard effect for hazard rates above the critical value.

Using a second-order Taylor expansion of f_{unc} around V_{eff} , we approximate the uncertainty multiplier in equation (5) as⁹

$$\frac{f_{unc}[V_0] - f_{unc}[V_1]}{[V_0 - V_1] f'_{unc}[V_{eff}]} \approx 1 + \underbrace{\frac{-f''_{unc}}{f'_{unc}} \bigg|_{V_{eff}} \left(V_{eff} - \frac{V_0 + V_1}{2} \right)}_{MHE_{contribution}^{unc}}$$

and obtain the overall impact from the marginal hazard effect

$$MHE^{total} \approx MHE^{neutral} [1 + MHE_{contribution}^{unc}].$$

The new term $MHE_{contribution}^{unc}$ is positive for hazards smaller than a critical value that is less than one half: the generalized mean V_{eff} places only the weight h on the bad outcome as compared to the arithmetic mean $\frac{V_0 + V_1}{2}$ that places equal weights on the two states of the world. Increasing uncertainty aversion increases the uncertainty contribution

⁸Recall that V_{eff} varies from V_0 to V_1 as the hazard goes from 0 to 1.

⁹The second order Taylor expansion for $f_{unc}(V_i)$ around V_{eff} delivers $f_{unc}(V_i) = f_{unc}(V_{eff}) + f'_{unc}(V_{eff})[V_i - V_{eff}] + \frac{1}{2}f''_{unc}(V_{eff})[V_i - V_{eff}]^2 + O([V_i - V_{eff}]^3)$ and results in the difference $f_{unc}(V_1) - f_{unc}(V_0) \approx f'_{unc}(V_{eff})[V_1 - V_0] + \frac{1}{2}f''_{unc}(V_{eff})([V_1 - V_{eff}]^2 - [V_0 - V_{eff}]^2)$. Inserting this difference into the left hand side of the equation and observing that $\frac{[V_0 - V_{eff}]^2 - [V_1 - V_{eff}]^2}{[V_0 - V_1]} = \frac{V_0^2 - V_1^2 + 2V_{eff}(V_0 - V_1)}{[V_0 - V_1]} = V_0 + V_1 + 2V_{eff}$ delivers the approximation on the right-hand side of the equation.

through the measure of absolute uncertainty aversion $\left. \frac{-f''_{unc}}{f_{unc}} \right|_{V_{eff}}$. However, it also reduces the contribution of the term in parentheses by reducing the uncertainty-averse mean V_{eff} . As a consequence, the marginal hazard effect is weakly concave in uncertainty aversion.

The uncertainty multiplier has a similar intuition as the expression for the willingness to pay for a risk reduction in the standard risk setting. A risk-averse agent is not always willing to pay more for a risk reduction than is a risk-neutral agent. The more risk-averse (uncertainty-averse) the decision-maker, the more he values wealth in bad states relative to good states. Giving up wealth for a hazard reduction makes a decision-maker even worse off in the bad state of the world. If a hazard is large enough, he prefers carrying wealth into the post-threshold world over spending it on reducing the hazard. This effect of uncertainty aversion on the MHE has parallels in the literature studying self-protection under uncertainty aversion (Treich 2010, Snow 2011, Alary et al. 2013, Maccheroni et al. 2013). In these settings, uncertainty aversion skews optimal policy towards reducing uncertainty, by making the less probable outcome even less likely. We see the same type of effect here.

3.2 The policy implications of the post-threshold world

In a world without uncertainty aversion, a potential tipping point also changes the policy trade-off through the terms proportional to the hazard rate in the first two lines of equation (3):

$$DWI^{neutral} \equiv h \left[\left(-\frac{\partial V_1}{\partial e_t} \right) - \left(-\frac{\partial V_0}{\partial e_t} \right) \right] \equiv h \left[\frac{\partial V_0}{\partial S_{t+1}} - \frac{\partial V_1}{\partial S_{t+1}} \right] \frac{\partial S_{t+1}(e_t^*)}{\partial e_t}. \quad (6)$$

Lemoine and Traeger (2014) call this adjustment the *differential welfare impact (DWI)* because it is proportional to the difference in the marginal impact of the control on the pre- and post-threshold value functions. This differential welfare impact strengthens policy if relaxing the control has a more harmful impact in the post-threshold regime. Note that this term does not depend on whether the probability of tipping is exogenous or endogenous.

We now analyze how uncertainty aversion changes the differential welfare impact. An uncertainty-averse policymaker is additionally averse to the types of poorly understood uncertainty that characterize tipping points. A strictly concave function f_{unc} captures her additional aversion to tipping point uncertainty. Our extended definition of DWI now collects all the changes in the first two lines of equation (3) with respect to the

continuation value:

$$DWI^{total} = DWI^{neutral} + \underbrace{h \left[\frac{f'_{unc}(V_1)}{f'_{unc}(V_{eff})} - 1 \right] \left(-\frac{\partial V_1}{\partial e_t} \right) + (1-h) \left[\frac{f'_{unc}(V_0)}{f'_{unc}(V_{eff})} - 1 \right] \left(-\frac{\partial V_0}{\partial e_t} \right)}_{DWI^{unc}}. \quad (7)$$

These changes capture the effect of an exogenous hazard and its interaction with uncertainty aversion. For an uncertainty-averse decision-maker who faces a welfare-decreasing tipping point, the concavity of f_{unc} implies $f'(V_0) < f'(V_{eff}) < f'(V_1)$. Therefore, the first term's contribution to DWI^{unc} in equation (7) is positive, whereas the second term's contribution is negative.

We now identify situations where DWI^{unc} is negative. We assume that the measure of absolute uncertainty aversion does not fall too quickly in welfare and that the tipping point does not cause too large a welfare loss. Then the following approximation characterizes the uncertainty contribution:¹⁰

$$DWI^{unc} \approx \frac{-f''_{unc}}{f'_{unc}} \Big|_{V_{eff}} \left[\underbrace{h(V_{eff} - V_1)}_{w^-} \left(-\frac{\partial V_1}{\partial e_t} \right) - \underbrace{(1-h)(V_0 - V_{eff})}_{w^+} \left(-\frac{\partial V_0}{\partial e_t} \right) \right]. \quad (8)$$

This contribution is proportional to the measure of absolute uncertainty aversion $\frac{-f''_{unc}}{f'_{unc}} \Big|_{V_{eff}}$ evaluated at V_{eff} . The terms in square brackets decide the sign of the contribution. Similarly to the uncertainty-neutral contribution $DWI^{neutral}$, this sign depends on the relative magnitudes of the control's marginal welfare impact in the pre- and post-threshold worlds. However, in contrast to the uncertainty-neutral contribution, these marginal welfare impacts are weighted by the terms w^+ and w^- , respectively.

Under uncertainty neutrality, we have $w^- = w^+ = h(1-h)(V_0 - V_1)$, so that each marginal welfare impact is weighted equally.¹¹ Introducing uncertainty aversion decreases V_{eff} . Assuming that the tipping point is harmful, uncertainty aversion thereby raises w^+ and reduces w^- . As a consequence, the uncertainty-averse contribution more strongly weights the marginal pre-threshold impact when determining the overall sign of the contribution. If this pre-threshold welfare impact of the control is larger than its post-threshold

¹⁰We approximate $f'_{unc}(V_i)$ for $i \in \{0, 1\}$ using the measure of absolute uncertainty aversion $\frac{-f''_{unc}}{f'_{unc}}$ evaluated at V_{eff} to obtain $\frac{f'_{unc}(V_i)}{f'_{unc}(V_{eff})} - 1 \approx \frac{f''_{unc}}{f'_{unc}} \Big|_{V_{eff}} (V_i - V_{eff}), i \in \{0, 1\}$.

¹¹To see this, recognize that $V_{eff} = (1-h)V_0 + hV_1$ under uncertainty neutrality. In this case, $V_{eff} - V_1 = (1-h)(V_0 - V_1)$ and $V_0 - V_{eff} = h(V_0 - V_1)$.

impact, then $DWI^{neutral} < 0$. Thus, equation (8) implies that, if the uncertainty-neutral contribution to the differential welfare impact is negative, then the uncertainty-averse contribution is also negative ($DWI^{unc} < 0$).

In general, the sign of the uncertainty-neutral contribution $DWI^{neutral}$ is ambiguous. In our climate change example, a tipping point that makes carbon dioxide more harmful in the post-threshold regime would tend to increase the optimal carbon tax through $DWI^{neutral}$. However, a tipping point with an endogenous hazard also tends to decrease the optimal carbon tax through $DWI^{neutral}$ because any reduction in carbon dioxide reduces the chance of triggering a tipping point in some future period. This reduction in future hazard rates is only valuable in the pre-threshold regime, prior to having triggered the tipping point.¹² If the marginal effect of emissions on the future tipping hazard is sufficiently painful, then the differential welfare impact turns negative: emission reductions are more valuable before crossing the threshold. Uncertainty aversion increases the policymaker's fear of tipping. First, this fear translates directly into the weights w^- and w^+ , which push DWI^{unc} towards becoming negative. Second, this fear of tipping raises the marginal pre-threshold benefit from emission reductions, which pushes both $DWI^{neutral}$ and DWI^{unc} towards becoming negative.

In our climate change application, the DWI increases the optimal carbon tax under uncertainty neutrality. The effect is driven by the stronger harm from emissions in the post-tipping world. Under uncertainty aversion, the DWI's sign switches because the policymaker becomes more afraid of tipping. This sign switch results from an increase in the pre-threshold value function's sensitivity to the tipping hazard. The MHE remains the primary channel through which tipping points affect the optimal tax on carbon dioxide emissions. It increases the cost of emissions because of the endogenous tipping hazard in the pre-threshold world, which the uncertainty-averse policymaker especially fears. If the system tips (at hazard rate h) the policymaker no longer benefits from a policy-induced reduction in the tipping hazard.

The effect of uncertainty aversion on the DWI has parallels in the literature studying self-insurance under uncertainty aversion (Gollier 2011, Traeger 2011, Alary et al. 2013). Our setting differs by considering the full intertemporal interaction between self-protection and self-insurance in a world with an endogenous tipping hazard. In particular, the cited

¹²The MHE captures the immediate chance of triggering a regime shift in a given period. However, a ton of carbon dioxide released in the current period also increases the chance of triggering a regime shift in future periods. The pre-threshold value function's response to an increase in carbon dioxide captures this increased chance of tipping in future periods. If there were more than one potential tipping point, then the post-threshold value function would also contain this additional payoff from reducing carbon dioxide emissions.

literature generally finds that self-insurance increases in uncertainty aversion.¹³ In our setting, uncertainty aversion increases the DWI only if the policy has a strong enough marginal welfare impact in the post-threshold world to dominate the increased fear of the endogenous tipping hazard. If triggering a tipping point were to end the world, then the marginal welfare impact in the post-threshold world would be zero. In this case, the DWI would reduce the stringency of policy, with or without uncertainty aversion (compare to Clarke and Reed 1994, Polasky et al. 2011).

4 A climate-economy model with tipping points and uncertainty aversion

We now consider the effect of climate tipping points on the optimal tax on carbon dioxide emissions. The optimal carbon tax equals the social cost of carbon when evaluated along the optimal policy path in a welfare-maximizing integrated assessment model. This section summarizes the model structure. We refer the reader to Lemoine and Traeger (2014) for the complete mathematical description and computational methods.

We reformulate the benchmark Dynamic Integrated model of Climate and the Economy (DICE) from Nordhaus (2008) as an infinite-horizon dynamic programming problem following Kelly and Kolstad (1999) and Traeger (2014). We introduce a tipping point in the climate system, optimal learning about the threshold that triggers a tipping point, and a generalized welfare evaluation. DICE is a Ramsey-Cass-Koopmans growth model that has an aggregate world economy interacting with a climate module. In each period, a policymaker allocates available output to consumption, to investment in a durable capital stock, and to reducing emissions of carbon dioxide (CO_2). Emissions increase the stock of CO_2 , which generates warming in future periods. Warming in turn reduces output. Reducing emissions therefore acts like investment in that it transfers wealth from the present to the future. The pre-threshold world has standard DICE dynamics along with the tipping possibility and temperature shocks calibrated to the historical record.

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. Current emissions therefore increase

¹³Note that, in contrast to the discussion at the end of Section 3.1, we are no longer talking about willingness to pay for a reduction in the hazard. In the context of a standard risk setting, we are here discussing how introducing risk changes the marginal value of savings when the marginal value of wealth varies by state of the world. In many settings, the assumption of prudence ($u'''(\cdot) > 0$) ensures that the marginal value of savings increases upon introducing a mean-zero risk to future income.

the chance of triggering a tipping point in some future period. 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. We evaluate two tipping points of concern in the climate science literature. In every model run, the policymaker faces a single tipping point and knows in advance what its effects would be. Lemoine and Traeger (2014) motivate each of the tipping points with reference to the scientific literature and describe the calibration of the endogenous hazard rate. They explore three strengths for each type of tipping point and discuss the implications of alternative modeling assumptions.

The first tipping point increases the climate feedbacks that amplify global warming, and the second increases the atmospheric lifetime of CO_2 . The first tipping point therefore increases the effect of emissions on temperature, and the second increases the time for which current emissions will affect the climate. Temperature dynamics in DICE depend on a parameter known as climate sensitivity, which is the equilibrium warming from doubling CO_2 . We represent a climate feedback tipping point as increasing climate sensitivity from its pre-tipping value of 3°C to either 4°C or 5°C . The second tipping point reflects the possibility that carbon sinks weaken beyond the predictions of coupled climate-carbon cycle models, which is similar to decreasing the “decay rate” of CO_2 . We represent these weakened sinks by decreasing the transfer of CO_2 out of the atmosphere by 25%, 50%, or 75%.

As noted in discussing equation (2), the function f_{unc} captures aversion to the Knightian uncertainty surrounding 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. 2009). Sticking to isoelastic preferences, we adopt the power function $f_{unc}(V) = ((1 - \eta)V)^{\frac{1-\gamma}{1-\eta}}$. Here γ is a measure of Arrow-Pratt relative risk aversion with respect to Knightian uncertainty, and $\eta = 2$ is the constant Arrow-Pratt measure of relative risk aversion used in the standard DICE utility function. The Knightian uncertainty surrounding tipping points contrasts with the standard risk posed by the historically-grounded temperature stochasticity governed by IP. If relative aversion to Knightian uncertainty γ coincides with standard relative risk aversion η , then the function f_{unc} is linear and drops out. In that case, the policymaker is uncertainty-neutral and the welfare evaluation is as in DICE. However, when $\gamma > \eta$, the function f_{unc} measures the policymaker’s additional aversion to the uncertainty (or ambiguity) governing tipping points as opposed to annual 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 aversion to tipping point uncertainty. The optimal carbon tax is the policy trajectory that maximizes the present value of net benefits within our extension of the benchmark DICE integrated assessment model. All of our graphs present results conditional on not having crossed the threshold: we want to understand how optimal policy changes in the face of a potential tipping point. The depicted paths draw the multiplicative temperature shock at its expected value in each period.

Figure 1 gives the effect of tipping points and uncertainty attitudes on the optimal carbon tax (social cost of carbon), the optimal CO₂ concentration path, and the optimal temperature path. All trajectories use the “middle” strength version of each tipping point. These plots vary γ from uncertainty-neutral ($\gamma = 2 = \eta$) to extremely uncertainty-averse ($\gamma = 100 > 2 = \eta$). We find that optimal policy is not highly sensitive to the policymaker’s level of uncertainty aversion. The near-term carbon tax varies by less than \$1/tCO₂ across the modeled range of γ , the peak CO₂ concentration varies by less than 30 ppm, and peak temperature varies by less than 0.2°C. For the climate feedback tipping point, the extreme form of uncertainty aversion with $\gamma = 100$ increases the optimal carbon tax by 6% in 2015 and by 12% in 2050. For the carbon sink tipping point, the extreme form of uncertainty aversion increases the optimal carbon tax by around 6.5% in both 2015 and 2050. The effect on near-term policy of varying uncertainty aversion across our scenarios is slightly smaller than the effect of varying the strength of a tipping point in Lemoine and Traeger (2014).

Figure 2 plots the year 2015 optimal carbon tax and the peak temperature for values of γ between 0 (uncertainty-loving) and 100 (strongly uncertainty-averse), with all graphs conditional on not having crossed the threshold. This figure depicts the results for three different strengths of the carbon sink tipping point and for two different strengths of the feedback tipping point.¹⁴ As the level of uncertainty aversion increases, the optimal year 2015 carbon tax becomes slightly stronger. The greatest effects arise in the case of the strongest carbon sink tipping point, with the plotted cases showing that the optimal tax increases from \$14/tCO₂ when $\gamma = 0$ to over \$16/tCO₂ when $\gamma = 100$. We also see that greater uncertainty aversion decreases optimal peak temperature slightly in all cases. The

¹⁴When the feedback tipping point increases climate sensitivity to 6°C, the large welfare losses associated with tipping cause computational problems in the presence of a high degree of uncertainty aversion.

Figure 1: Time paths for the optimal carbon tax (current value), the CO₂ 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 potential tipping point over time. The parameter γ controls the strength of uncertainty aversion, ranging from uncertainty neutrality ($\gamma = 2$) to extreme uncertainty aversion ($\gamma = 100$).

Figure 2: The optimal carbon tax in 2015 and the peak temperature reached for each type of uncertainty attitude. The plotted simulations assume expected draws of the temperature shock and assume that the tipping point happens to never occur. The parameter γ controls the strength of uncertainty aversion, ranging from slightly uncertainty-loving ($\gamma = 0$) to uncertainty-neutral ($\gamma = 2$) to extremely uncertainty-averse ($\gamma = 100$).

greatest effects on peak temperature arise for stronger versions of the feedback tipping point, with the plotted cases showing that peak temperature falls from 2.9°C in the case of $\gamma = 0$ to 2.7°C in the case of $\gamma = 100$.

We now use our decomposition of the optimal emission tax from Section 3 to assess the channels through which uncertainty aversion affects policy (Figure 3). Each panel depicts either the differential welfare impact (DWI) or the marginal hazard effect (MHE) in the case of uncertainty neutrality ($\gamma = 2$) and in the case of extreme uncertainty aversion ($\gamma = 100$). Begin with the MHE (top row). In keeping with our analysis, we see that the total MHE is positive in all cases and that uncertainty aversion can substantially increase the MHE. Further, the MHE rises strongly over time as the welfare loss from tipping increases and as the marginal effect of emissions on the hazard increases through learning about the threshold's location (see Lemoine and Traeger 2014). The MHE grows faster in the face of the feedback tipping point than in the face of the carbon sink tipping

point because the welfare loss imposed by the feedback tipping point increases strongly in CO_2 . Once the feedback tipping point has occurred, high CO_2 concentrations cause much greater damages than in the pre-threshold regime. In contrast, the welfare loss from the carbon sink tipping point starts out relatively high, but it does not grow as fast over time because the policymaker's ability to control the cost of tipping is not as state-dependent.¹⁵

Now consider the DWI (middle row). It is typically more than an order of magnitude smaller than the MHE. In the case of uncertainty neutrality, the DWI is mostly positive because emission reductions pay off especially strongly in the tipped regimes, where warming responds more strongly to CO_2 concentrations or the effectiveness of carbon sinks has been substantially reduced. The exception arises for the carbon sink tipping point late in the century. There, we see a slightly negative DWI because the tipping point makes carbon dioxide reach very high levels, which reduces the marginal effect of another unit of carbon dioxide on temperature.¹⁶

Under strong uncertainty aversion, the total DWI reduces the optimal carbon tax, through both its uncertainty-averse and uncertainty-neutral components. Section 3.2 explains how uncertainty aversion can switch the sign of the DWI by making the pre-threshold value function more sensitive to the control than is the post-threshold value function. In particular, imagine that pre-threshold welfare is strongly affected by later periods' tipping hazards. Uncertainty aversion amplifies the policymaker's fear of future tipping, which remains a hazard as long as she is still in the pre-tipping world and temperatures are increasing. This impact is measured by future periods' MHE. Increasing future periods' MHE works to increase the marginal effect of emission reductions in the pre-threshold world (i.e., to increase $-\partial V_0/\partial e_t$). The differential welfare impact corrects for the fact that, with probability h , the policymaker will no longer be in this pre-threshold world after today. If the policymaker does happen to enter the post-threshold world, then today's abatement no longer provides the additional payoff of reducing future tipping hazards. Fearing future tipping hazards therefore does not change $\partial V_1/\partial e_t$. If the

¹⁵See the discussion in Lemoine and Traeger (2014) for more on how the two types of tipping points affect policy in qualitatively different ways.

¹⁶Additional units of carbon dioxide generate warming by trapping outgoing infrared radiation (heat). When there is already a lot of CO_2 in the atmosphere, the stock of CO_2 nearly completely blocks the wavelengths over which CO_2 most effectively traps heat. Additional emissions therefore primarily contribute to warming by blocking wavelengths at which they are less effective, so that the marginal contribution of emissions to warming declines in the stock of CO_2 . The negative DWI tells us that the marginal impact of emissions in the post-threshold world no longer dominates its impact in the pre-threshold world, which includes both the immediate impact on welfare in the pre-threshold world as well as the reduction in the chance of future tipping, captured in the pre-threshold value function's response to emissions.

Figure 3: Time paths for MHE (top) and DWI (middle) for tipping points that increase climate sensitivity to 5°C (left) and that weaken carbon sinks by 50% (right). The bottom row shows DWI calculated using a “no-threshold” continuation value for the pre-tipping regime, as discussed in the text. The lines labeled “Unc” depict the contribution of uncertainty aversion in cases with $\gamma = 100$. Note that the scales of the vertical axes do not match each other.

(pre-threshold) fear of tipping is sufficiently large (making $-\partial V_0/\partial e_t$ large), then it makes emission reductions more valuable in the pre-threshold world ($-\partial V_0/\partial e_t > -\partial V_1/\partial e_t$), so that accounting for the differential welfare impact reduces the optimal carbon tax ($DWI < 0$).

The bottom row of Figure 3 demonstrates this reasoning. Instead of using the correct pre-threshold continuation value $V_0(S_{t+1})$, it calculates a (hypothetical) DWI using the continuation value that would obtain if the policymaker were not aware that current emissions increased the chance of triggering tipping points in future periods. In this case, each period's policymaker believes that the only hazard is that tipping might occur during the current period. This calculation removes the chance of triggering a future tipping point from the pre-threshold continuation value. Now emission reductions are unambiguously more valuable if the policymaker happens to cross a threshold. In contrast to the previous paragraph's discussion, it is unambiguous that $-\partial V_0/\partial e_t < -\partial V_1/\partial e_t$. The bottom row of Figure 3 indeed shows that, in this hypothetical world, both the uncertainty-neutral and the uncertainty-averse DWI contributions are positive: when a policymaker only considers a one-time chance of tipping, each ton of carbon dioxide more strongly impacts welfare in the post-threshold word.

6 Conclusions

We have shown how aversion to uncertainty (or ambiguity) affects optimal policy in the face of poorly understood tipping points. First, uncertainty aversion affects optimal policy through the marginal hazard effect (MHE) when the probability of triggering a tipping point is endogenous. The MHE captures how introducing the chance of a harmful tipping point moves optimal policy in a direction that reduces the probability of tipping. Uncertainty aversion tends to move policy in a direction that reduces Knightian uncertainty. More precisely, for a tipping hazard that is not too large, uncertainty aversion pushes the policymaker to spend more on reducing the endogenous hazard. If the hazard becomes large enough, the value of transferring consumption into the bad state of the world can outweigh the self-protection incentive and reduce spending on avoiding the hazard. This effect of uncertainty aversion increases in the sensitivity of the tipping hazard to the policy control, in the expected welfare loss from tipping, and, at least for small hazard rates, in uncertainty aversion. These findings relate closely to the usual arguments about self-protection under risk and ambiguity.

Second, uncertainty aversion changes the differential welfare impact (DWI) of the

optimal policy. The DWI captures how the marginal value of a control depends on whether a tipping point is triggered or not. If the payoffs from a stringent policy are higher in the post-threshold regime than in the pre-threshold regime, then the DWI strengthens the optimal policy if the chance of future tipping is not too great and the policymaker is not too uncertainty-averse. However, strong aversion to an endogenous tipping hazard can flip this sign. The fear of eventual tipping makes a strong policy particularly valuable in the pre-threshold (and only in the pre-threshold) world. Then the post-threshold impact adjustment captured by the DWI reduces the strength of the optimal policy.

In our numeric climate change application, we find that uncertainty aversion slightly increases the optimal carbon tax. The DWI is slightly positive under uncertainty neutrality, but it reduces the optimal carbon tax in the case of an uncertainty-averse policymaker. Uncertainty aversion's effect on the MHE always dominates its effect on the DWI. The carbon tax increase remains moderately small even for high degrees of ambiguity aversion. Considering stronger tipping points would increase both the uncertainty-neutral and the uncertainty-averse contributions to the optimal carbon tax. Increasing the tipping hazard induced by a given temperature step can increase the uncertainty-averse contribution, until the tipping hazard grows too large and the contribution falls again. Tipping point must be very severe if uncertainty aversion is to substantially increase the optimal carbon tax relative to other contributions.

In reality, there is a severe lack of knowledge about the existence and location of tipping points in the climate system. We find that a framework that explicitly incorporates the Knightian nature of this uncertainty leads to only small (upwards) adjustments of the optimal carbon tax. Jensen and Traeger (2014) find similarly small adjustments from introducing uncertainty aversion in a world of smooth climate change. Further, Lemoine and Traeger (2014) show that a better estimate of the Bayesian prior can have a stronger impact on the optimal policy than the adjustments we obtain here from accounting for (or eliminating) the Knightian nature of the uncertainty. Our results suggest that modelers and policymakers should not shy away from following a best-guess Bayesian prior approach to tipping points when deriving optimal mitigation policy.

References

- Alary, David, Christian Gollier, and Nicolas Treich (2013) "The effect of ambiguity aversion on insurance and self-protection," *The Economic Journal*, Vol. 123, No. 573, pp. 1188–1202.

- 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*, Vol. 299, No. 5615, pp. 2005–2010.
- Anderson, Evan, William Brock, Lars Peter Hansen, and Alan H. Sanstad (2014) "Robust analytical and computational explorations of coupled economic-climate models with carbon-climate response," RDCEP Working Paper 13-05, The Center for Robust Decision Making on Climate and Energy Policy.
- Bansal, Ravi and Amir Yaron (2004) "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, Vol. 59, No. 4, pp. 1481–1509.
- Barrett, Scott and Astrid Dannenberg (2012) "Climate negotiations under scientific uncertainty," *Proceedings of the National Academy of Sciences*, Vol. 109, No. 43, pp. 17372–17376.
- (2014) "Sensitivity of collective action to uncertainty about climate tipping points," *Nature Climate Change*, Vol. 4.
- Berger, Loïc, Johannes Emmerling, and Massimo Tavoni (2015) "Managing catastrophic climate risks under model uncertainty aversion," *Management Science*, Vol. forthcoming.
- Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek (2013) "The social cost of stochastic and irreversible climate change," Working Paper 18704, National Bureau of Economic Research.
- Camerer, Colin and Martin Weber (1992) "Recent developments in modeling preferences: Uncertainty and ambiguity," *Journal of Risk and Uncertainty*, Vol. 5, No. 4, pp. 325–370.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011) "Uncertainty averse preferences," *Journal of Economic Theory*, Vol. 146, No. 4, pp. 1275–1330.
- Clarke, Harry R. and William J. Reed (1994) "Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse," *Journal of Economic Dynamics and Control*, Vol. 18, No. 5, pp. 991–1010.
- 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.

- Diekert, Florian K. (2015) “Threatening thresholds? The effect of disastrous regime shifts on the cooperative and non-cooperative use of environmental goods and services,” Memorandum 12/2015, Oslo University, Department of Economics.
- 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.
- Fitzpatrick, Luke G. and David Kelly (2015) “Probabilistic Stabilization Targets,” *Working Paper*.
- Gilboa, Itzhak and Massimo Marinacci (2013) “Ambiguity and the Bayesian paradigm,” in Daron Acemoglu, Manuel Arellano, and Eddie Dekel eds. *Advances in Economics and Econometrics: Tenth World Congress, Volume I, Economic Theory*, New York, NY: Cambridge University Press, pp. 179–242.
- Gjerde, Jon, Sverre Grepperud, and Snorre Kverndokk (1999) “Optimal climate policy under the possibility of a catastrophe,” *Resource and Energy Economics*, Vol. 21, No. 3-4, pp. 289–317.
- Gollier, Christian (2011) “Portfolio choices and asset prices: The comparative statics of ambiguity aversion,” *The Review of Economic Studies*, Vol. 78, No. 4, pp. 1329–1344.
- Heal, Geoffrey (1984) “Interactions between economy and climate: A framework for policy design under uncertainty,” in V. Kerry Smith and Ann Dryden White eds. *Advances in Applied Microeconomics*, Vol. 3, Greenwich, CT: JAI Press, pp. 151–168.
- Hennlock, Magnus (2009) “Robust control in global warming management: An analytical dynamic integrated assessment,” discussion Paper 09-19, Resources For the Future.
- Hwang, In Chang, Frédéric Reynès, and Richard S. J. Tol (2014) “The effect of learning on climate policy under fat-tailed uncertainty,” Munich Person RePEc Archive Paper 53681.
- Jensen, Sverre and Christian Traeger (2013) “Optimally climate sensitive policy: A comprehensive evaluation of uncertainty and learning,” working Paper, Department of Agricultural and Resource Economics, University of California Berkeley.
- Jensen, Sverre 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.
- Keller, Klaus, Benjamin M. Bolker, and David F. Bradford (2004) “Uncertain climate

- thresholds and optimal economic growth,” *Journal of Environmental Economics and Management*, Vol. 48, No. 1, pp. 723–741.
- Kelly, David L. and Charles D. Kolstad (1999) “Bayesian learning, growth, and pollution,” *Journal of Economic Dynamics and Control*, Vol. 23, No. 4, pp. 491–518.
- Kelly, David L. and Zhuo Tan (2015) “Learning and climate feedbacks: Optimal climate insurance and fat tails,” *Journal of Environmental Economics and Management*, Vol. 72, pp. 98–122.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005) “A smooth model of decision making under ambiguity,” *Econometrica*, Vol. 73, No. 6, pp. 1849–1892.
- (2009) “Recursive smooth ambiguity preferences,” *Journal of Economic Theory*, Vol. 144, No. 3, pp. 930–976.
- 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*, Vol. 106, No. 13, pp. 5041–5046.
- Lange, Andreas and Nicolas Treich (2008) “Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions,” *Climatic Change*, Vol. 89, No. 1, pp. 7–21.
- Leach, Andrew J. (2007) “The climate change learning curve,” *Journal of Economic Dynamics and Control*, Vol. 31, No. 5, pp. 1728–1752.
- 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.
- Lenton, Timothy M., Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber (2008) “Tipping elements in the Earth’s climate system,” *Proceedings of the National Academy of Sciences*, Vol. 105, No. 6, pp. 1786–1793.
- Li, Xin, Borghan Narajabad, and Ted Temzelides (2014) “Robust Dynamic Optimal Taxation and Environmental Externalities,” *CESifo Working Paper Series*, Vol. 4562.
- 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.

- Maccheroni, Fabio, Massimo Marinacci, and Doriana Ruffino (2013) “Alpha as ambiguity: Robust mean-variance portfolio analysis,” *Econometrica*, Vol. 81, No. 3, pp. 1075–1113.
- Millner, Antony, Simon Dietz, and Geoffrey Heal (2013) “Scientific ambiguity and climate policy,” *Environmental and Resource Economics*, Vol. 55, No. 1, pp. 21–46.
- Nævdal, Eric (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*, Vol. 30, No. 7, pp. 1131–1158.
- Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa (2013) “Crises and recoveries in an empirical model of consumption disasters,” *American Economic Journal: Macroeconomics*, Vol. 5, No. 3, pp. 35–74.
- Nordhaus, William D. (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven: Yale University Press.
- van der Ploeg, Frederick (2014) “Abrupt positive feedback and the social cost of carbon,” *European Economic Review*, Vol. 67, pp. 28–41.
- Polasky, Stephen, Aart de Zeeuw, and Florian Wagener (2011) “Optimal management with potential regime shifts,” *Journal of Environmental Economics and Management*, Vol. 62, No. 2, pp. 229–240.
- 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*, Vol. 105, No. 38, pp. 14245–14250.
- Rudik, Ivan (2015) “The fragility of robustness: Climate policy when damages are unknown,” working Paper, Iowa State University.
- 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.-M. Fussel, A. B. Pittock, A. Rahman, A. Suarez, and J.-P. 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*, Vol. 106, No. 11, pp. 4133–4137.
- Snow, Arthur (2011) “Ambiguity aversion and the propensities for self-insurance and self-protection,” *Journal of Risk and Uncertainty*, Vol. 42, No. 1, pp. 27–43.
- 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. (2010) “Subjective risk, confidence, and ambiguity,” CUDARE Working Paper 1103, University of California, Berkeley.
- (2011) “Discounting and Confidence,” *CUDARE working paper 1117*.
- (2014) “Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity,” *Economic Theory*, Vol. 56, No. 3, pp. 627–664.
- Treich, Nicolas (2010) “The value of a statistical life under ambiguity aversion,” *Journal of Environmental Economics and Management*, Vol. 59, No. 1, pp. 15–26.
- Tsur, Yacov and Amos Zemel (1996) “Accounting for global warming risks: Resource management under event uncertainty,” *Journal of Economic Dynamics and Control*, Vol. 20, No. 6-7, pp. 1289–1305.
- (2009) “Endogenous discounting and climate policy,” *Environmental and Resource Economics*, Vol. 44, No. 4, pp. 507–520.
- Vissing-Jørgensen, Annette and Orazio P. Attanasio (2003) “Stock-market participation, intertemporal substitution, and risk-aversion,” *The American Economic Review*, Vol. 93, No. 2, pp. 383–391.
- Weil, Philippe (1989) “The equity premium puzzle and the risk-free rate puzzle,” *Journal of Monetary Economics*, Vol. 24, No. 3, pp. 401–421.
- de Zeeuw, Aart and Amos Zemel (2012) “Regime shifts and uncertainty in pollution control,” *Journal of Economic Dynamics and Control*, Vol. 36, No. 7, pp. 939–950.