DICESC: Optimal Policy in a Stochastic Control Framework

Charles W. Chang

charles.w.chang@aae.wisc.edu

Optimization Group, Wisconsin Institute for Discovery
Department of Agricultural and Applied Economics, University of Wisconsin-Madison

Selected Paper prepared for presentation at the Agricultural & Applied Economics Association's 2014 AAEA Annual Meeting, Minneapolis, MN, July 27-29, 2014.

DICESC: Optimal Policy in a Stochastic Control Framework

Charles W. Chang
charles.w.chang@aae.wisc.edu [1]

[1]University of Wisconsin-Madison

June 19, 2014

Abstract

The effects of anthropogenic greenhouse gas emission to climate change are uncertain; how does this uncertainty affect the optimal carbon abatement policy? Tackling this classic problem is challenging; not only does it require integrated assessment analyses that consolidate an adequate representation of both climate and economy, but also a sensible understanding of hedging incentives and decision making under uncertainty. In the Integrated Assessment Model (IAM) community, this representation is severely restricted by computational burdens and model intractability (Webster, Santen, and Parpas, 2012). I provide a minimal framework that relies on the transparency of mathematical programming, to assess precautionary mitigation as a means to hedge against the risk and uncertainty of climate catastrophes. The scope of the paper encompasses the following:

- assessment of optimal abatement under parameter uncertainty and endogenous learning;
- decomposition of optimal abatement with respect to mitigatory incentives;
- assessment of policy sensitivity to assumed distributions of uncertain temperature thresholds.

I introduce DICESC, a multistage stochastic control version of Nordhaus' DICE model to demonstrate that policy optimization in which the hazard rates of climate catastrophes can be controlled significantly increases incentives for precautionary abatement. Despite the drastic level increase in precautionary abatement the proposed formulation achieves, I stress the importance of understanding the distributive properties of uncertainty for the robust assessment of optimal policy.

1. Introduction

An increase in global atmospheric concentration of carbon dioxide (CO_2) and other industrial greenhouse gases (GHGs) such as methane and chlorofluorocarbons (CFCs), will result in climate damages. Although the full damage impacts of carbon emissions is uncertain, the series of reports by the Intergovernmental Panel on Climate Change (IPCC) conclude that the accumulation of GHGs bring about higher global surface temperature, greater variability in precipitation levels, rise in sea levels and more frequent extreme climate catastrophes (Metz, Davidson, Bosch, Dave, and Meyer, 2007).

In addressing the economic impacts of climate change, numerical Integrated Assessment Models (IAMs) of varying complexity have assumed a significant role; i.e. an integrated framework that models climate science (how GHGs function to increase atmospheric temperature), the projected abatement costs and climate damages to economic welfare (Pindyck, 2013). Well known IAMs that utilize optimal economic growth models to couple economy and climate such as DICE (Nordhaus and Sztorc, 2013; Nordhaus, 1994) and MERGE (Manne, Mendelsohn, and Richels, 1995; Manne and Richels, 1992), optimize the trade off between abatement today and climate damages in an uncertain future. (Keller, Bolker, and Bradford, 2004).

Despite the growing complexity of climate IAMs, Pindyck (2013) and others (Kaufman, 2012; Kelly and Kolstad, 1999) point out that the elements of uncertainty in climate change policy can present seriously misleading optimal abatement policies. Pindyck (2013) further argues that the most important driver of the social cost of carbon is the possibility of a catastrophic climate outcome; the likelihood of a catastrophic impact triggered by an increase in atmospheric temperature past a temperature threshold. Hence, a scenario of active precautionary abatement is non-credible without considering the possibility of climate induced catastrophes. The need for incorporating catastrophes to climate IAMs was acknowledged previous to both Pindyck and Kaufman; policy-relevant climate induced catastrophes that are assumed to take place when atmospheric temperature increase exceeds a *tipping point*, has entered numerous IAM models.

Representing catastrophes such as the Atlantic Thermohaline Circulation collapse (THC) and the Greenland ice sheet melting to optimal growth models is technically challenging; optimal climate policy optimizes between climate damages and lost opportunities for economic growth and must also consider the non-linearity of tipping elements and hysteresis. To make matters worse, the combined effects of uncertainty in the global climate system and non-linearity of tipping elements can pose intractable computational problems in policy assessment (Keller, Bolker, and Bradford, 2004). Including Keller et al.'s work, a number of studies have integrated a stochastic framework into IAMs such as DICE to demonstrate the policy response to low probabilities and high consequences of catastrophic outcomes; major studies include Manne and Richels (1992), Ackerman, Stanton, and Bueno (2010), Kelly and Kolstad (1999) , Cai, Judd, and Lontzek (2012), Hope (2006), Lemoine and Traeger (2010), Lemoine and Traeger (2013), Keller, Bolker, and Bradford (2004), Webster (2000).

Recent work shows that the scientific understanding of the climate system has improved and awareness of the relevance of uncertainty modeling has heightened. As a result climate models have become larger and more complex, yet the uncertainty of the climate system still remains predominantly unresolved (Urban, Holden, Edwards, Sriver, and Keller, 2014; Pindyck, 2013). Efforts to formulate an intricate decision analysis framework in response to an uncertain climate outcome, outlining a lower bound to near-term abatement is consequently deemed important given the slow rates of climate learning that takes place. I review the characteristics of stochastic decision analysis frameworks currently employed by IAMs.

On one hand, efforts to formulate the optimal policy in response to climate uncertainty have conveniently steered the IAM community to integrate recursive dynamic programming frameworks to economy-climate IAMs. However obtaining a numerical solution through dynamic programming has proved to be computationally burdensome and dimensionally restricted, especially for assessing sequential decision policies characterized by a large set of *actions* (Dixit, 1990).

On the other hand, efforts to adequately represent the structure of uncertainties and the non-linear character of climate damages introduce computational problems such as non-smooth gradients and local optima (Keller, Bolker, and Bradford, 2004). Such cases may limit the use of commercially available optimization solvers to obtain numerical model solutions. Detailed representations of the integrated system can moreover restrict the embedding of stochasticity within the model framework itself and confine uncertainty analyses to Monte Carlo simulations (Webster, Santen, and Parpas, 2012).

The difficulties IAMs face point out the need for transparent frameworks that allow for sufficient expansion to the scope of mitigatory actions, while maintaining model tractability. In this paper, I present DICESC, a multistage stochastic programming with recourse formulation of Nordhaus' DICE2013 in which the climate hazard rates are endogenously controlled. The model is conveniently written as a mathematical program for the framework's action set to be easily expanded to hedge against uncertain risky outcomes. More importantly, DICESC maintains a recourse structure that avoids the scaling limitations of classical multistage stochastic programming, and can be solved in a tractable way using commercial optimization solvers (Defourny, Ernst, and Wehenkel, 2012).

Previous Studies: Stochastic Framework for Modeling Tipping Points

Unlike IAMs that focus on the avoidance of market and non-market damages as an incentive for near-term mitigation, recent work by Lemoine and Traeger (2013) and Cai, Judd, and Lontzek (2012), along with previous work of Keller, Bolker, and Bradford (2004) have introduced a precautionary motive for mitigation by modeling the uncertainty of climate tipping points. The three papers aforementioned all incorporate stochastic irreversible climate damages to Nordhaus' DICE model and contribute uniquely to the stochastic formulations of assessing the impacts of uncertain climate thresholds.

The work of Lemoine and Traeger (2013) and Cai, Judd, and Lontzek (2012) use recursive dynamic programming methods in their model formulations. Although both model irreversible climate disasters that impose a permanent shock to economic production, each paper employs contrasting assumptions to characterize uncertainty and risky outcomes. Lemoine and Traeger (2013) use a continuous distribution of tipping points for which endogenous learning about the tipping points reduces the uncertainty domain of threshold temperatures, whereas Cai et al.'s DSICE uses discrete markov states with no scope of learning. Nevertheless both conclude that uncertain possibilities of tipping climate increases the overall level of abatement, and further demonstrate that near-term mitigation is desirable since it reduces the rate of temperature change, thereby decreasing the likelihood of catastrophic impacts.

In more detail, DSICE assumes that a climate catastrophe that permanently reduces economic output by 10 percent occurs when atmospheric temperatures trigger a tipping element. As a result, the optimal abatement trajectory increases monotonically until the tipping point and drastically decreases thereafter, representing the economy's anti-tipping efforts despite the rising costs of abatement. In the stochastic setting, the probability of transitioning to a markov state characterized by a lower tipping point (worse state) is a function of the atmospheric temperature at each time period; higher temperatures lead to higher probability of moving to a worse state. Lemoine and Traeger (2013) additionally establishes the role of Bayesian learning about the location of the threshold. The policy maker in this model acts and observes the climate system's response; at each time period the economy realizes if its actions have triggered climate tipping. Through learning, the economy updates its prior belief in the tipping point distribution and updates the effective hazard rate of crossing the threshold at each time.

Keller, Bolker, and Bradford (2004)'s model, *FRANC* in contrast, puts more emphasis on the non-linearity of hysteresis modeling and models the risk of climate tipping by introducing uncertainty to the climate sensitivity parameter. In other words, uncertainty of tipping lies not in temperature, but in climate sensitivity. In contrast to Lemoine and Traeger's representation, tipping points in this formulation denote the *critical equivalent carbon dioxide concentration,* P_{CO_2Crit} assumed for each climate sensitivity parameter value included in the uncertainty domain. Characterization of uncertainty in climate sensitivity as well as threshold damages are calibrated to the IPCC's AR4 (Metz, Davidson, Bosch, Dave, and Meyer, 2007), to output refined non-convex abatement trajectories for each deterministic state of the world. Keller et al. shows that optimal policy that weighs the costs and total discounted utility of catastrophe mitigation at every period, can in cases of very high climate sensitivity, choose a less stringent abatement path than in the case of lower sensitivity. The stochastic setting applies a Monte Carlo simulation to randomly draw the parameters of uncertainty, and does not provide room for endogenous hazard rates. Lastly, expected utility is formulated as a single policy that maximizes the sum of utilities for the sampled deterministic states.

Keller, Bolker, and Bradford (2004); Lemoine and Traeger (2013); Cai, Judd, and Lontzek (2012) all demonstrate optimal hedging strategies under an *act-then learn* decision problem; the policy maker makes decisions for the future under uncertainty of the tipping element, which motivates the optimal hedging strategy to adopt abatement levels that lie among all stochastic states (Manne and Richels, 1992). Near-term abatement is then justified by adverse stochastic states that experience catastrophic disasters and increase the expected abatement trajectory. In addition to the stochastic nature of the decision problem, endogenous hazard rates prove to further drive precautionary abatement by incentivizing the economy to decrease the likelihood of climate tipping.

For improved representation of the climate system however, the majority of policy assessments using IAMs are deterministic. They are characterized by a *learn-then act* framework in which policy optimization takes place *given* perfect information of the state. I point out that evaluating the influence of parameter uncertainty to optimal policy in a *learn-then act* framework cannot provide a fully comprehensive assessment and further demonstrate how the framework characterization can result in seriously misleading policies.

In the following section, I provide a minimalist framework of precautionary mitigation using a *stochastic control* version of DICE2013, DICESC, to investigate the impact of uncertain catastrophic loss on near term mitigation. Following Lemoine and Traeger 2013, I adopt endogenous learning of the tipping point distribution, which accordingly updates the hazard rate at each time period. *Stochastic control* in this context, refers to stochastic programming with recourse in which the probabilities of each state of the world taking place is a function of the decision variables today. For robustness of the optimal policy, the model builds upon the stochastic control framework assuming only the distribution of tipping points and the Bayesian learning mechanism through which the hazard rate is determined. I further stress the computational transparency of the model as DICESC is formulated as a one-shot optimization problem that can replicate results from the relatively complex recursive methods presented in Lemoine and Traeger (2013) and Cai, Judd, and Lontzek (2012), in a matter of minutes.

2. DICESC: The Model

2.1 Introduction to DICE

DICESC uses the 2013 version of the DICE model developed by Nordhaus to demonstrate the relevance of stochastic control. DICE2013 (hereafter cited as DICE) is an updated version of the original climate IAM authored in 1994 (Nordhaus, 1994), calibrated to be consistent with the IPCC Fifth Assessment Report (Stocker, Qin, Plattner, Tignor, Allen, Boschung, Nauels, Xia, Bex, Midgley, et al., 2013). As one of the first policy optimizing climate IAMs, DICE describes the impact of GHG emissions control to the economy which is governed by the model's carbon cycle and climate system. In this section, I briefly state the main equations of interest in each sector for readers who are unacquainted with the model.

$$W = \sum_{t=1}^{Tmax} \frac{1}{(1+\rho)^t} U[c(t), L(t)] \tag{1}$$

$$U[c(t), L(t)] = L(t) \left[\frac{c(t)^{1-\alpha}}{(1-\alpha)} \right] \tag{2}$$

$$E_{Ind}(t) = \sigma(t)[1-\mu(t)]Q(t) \tag{3}$$

W in equation (1) is the social welfare function, where $U(c(t), L(t))$ denotes the utility of consumption in period t. It is important to note that $c(t)$ represents the decision variable, per-capita income, whereas $L(t)$ is an exogenous parameter that represents population size. The objective of the model is to maximize the sum of the discounted value of utility through the economic horizon (years 2015 to 2115). Also note that ρ is the pure rate of time preference and α, the elasticity of marginal utility of consumption, which are parameterized to be 0.15 and 1.45, respectively. Equation (3) describes the amount of GHG emitted through economic production, $Q(t)$; here $\sigma(t)$ denotes the GHG emissions to output ratio and $\mu(t)$, emissions abatement or the fractional reduction in carbon emissions. While the emissions output ratio is an exogenous parameter that decreases every period at a fixed rate, $\mu(t)$ is the decision variable central to this research.

$$Q(t) = \frac{[1-\Lambda(t)]A(t)K(t)^{\gamma}L(t)^{1-\gamma}}{[1+\Omega(t)]} \tag{4}$$

$$\Omega(t) = \Psi_1 T_{AT}(t) + \Psi_1[T_{AT}(t)]^2 + d_{cat} \cdot 2 \left[.00644 \left(\frac{T_{AT}(t)}{\bar{T}_{AT}} \right)^3 \right] \tag{5}$$

$$\Lambda(t) = \theta_1(t)\mu(t)^{\theta_2} \tag{6}$$

Equations (4)-(6) show the dynamics of climate damages and costs taking place as a fraction of gross output. Q denotes economic output net of climate damages ($\Omega(t)$) and abatement costs ($\Lambda(t)$), where $A(t)$ is the Hicks-neutral total factor productivity and $K(t)$, the capital stock. Equation (6) expresses the ratio of abatement costs to output, ($\Lambda(t)$), as a convex function of abatement and (5) presents damage impacts from climate change as a function of atmospheric temperature. The third term in equation (5) represents catastrophic damages, which drastically increase climate damages when atmospheric temperature moves past a threshold temperature, $\bar{T}_{AT}$. The dummy parameter d_{cat} conveniently enables the modeling of *irreversible* climate catastrophes.

$$M_{AT}(t) = E_{Ind}(t) + E_{Land}(t) + \phi_{11}M_{AT}(t-1) + \phi_{21}M_{UP}(t-1) \tag{7}$$

$$F(t) = \eta log_2 \left[\frac{M_{AT}(t)}{M_{AT}(1750)}\right] + F_{EX}(t) \tag{8}$$

$$T_{AT}(t) = T_{AT}(t-1) + \xi_1 F(t) - \xi_2 T_{AT}(t-1) - \xi_3[T_{AT}(t-1) - T_{LO}(t-1)] \tag{9}$$

The last set of equations describe the process of how industrial emissions increase the mean atmospheric temperature. Equation (7) shows that emissions adds to the carbon concentration in the atmosphere $M_{AT}(t)$, which in turn increases the radiative forcing of greenhouse gases $F(t)$ (8). Lastly and most importantly, equation (9) displays the response of $T_{AT}(t)$, the mean atmospheric (surface) temperature, to the increase in radiative forcing.

Parts of DICE are deemed as being overly simplified and to lack empirical support; to settle on a concrete policy that requires minimal margin for error such as the social cost of carbon, a much more detailed IAM should be used (Nordhaus, 2011). However, the transparency of DICE is convenient to introduce methodological steps that can suggest effective directions in which optimal policy should move toward, a reason why the model is discussed so extensively for decision analysis in climate risk and uncertainty. The small size of the model also allows for policy optimization problems to be solved using conventional optimization tools.

2.2 The Model

This section presents specific modifications made to DICE. While maintaining the simple non-linear optimization format of DICE, I introduce a recourse tree that can integrate various states of the world.

Recourse

Note again that the time horizon t of economic activity extends to 2115 with time periods of five years ($t = 2015, 2020, \ldots$), while the time horizon of climate evolution and damages extends to year 2300. To this structure I add *states of the world*, s; scenarios each associated with the year in which catastrophic damage takes place (see *Figure 1.*). *Figure 1.* gives the planner (who at the moment lives in $t = 2015$) a recourse tree of four states of the worlds (SoWs); $s = 2015$ represents the unlikely scenario that an irreversible climate catastrophe occurs at year 2015, which the planner will only realize in the subsequent period, year 2020.

The *never* state represents a world in which catastrophic outcomes do not take place and corresponds to the recourse path $[1 \rightarrow 2 \rightarrow 4 \rightarrow 7]$. Then for SoW 2020 and 2025, the corresponding paths initially share the *never* path but deviate down the catastrophe path at time 2020 and 2025, respectively. In each SoW, the d_{cat} in equation (5) takes the value 1 at

$t = s$, to model an irreversible catastrophic outcome that hits at the corresponding time. It is important to note that with recourse, a SoW varying temperature threshold is applied (along with other SoW varying decision variables) to equation (5) resulting in a state varying damage equation (10).

$$\Omega(s,t) = \Psi_1 T_{AT}(s,t) + \Psi_1[T_{AT}(s,t)]^2 + d_{cat} \cdot 2\left[.00644\left(\frac{T_{AT}(s,t)}{\bar{T}_{AT}(s)}\right)^3\right] \tag{10}$$

$$\bar{T}_{AT}(s) = T_{AT}(never, t = s) \tag{11}$$

I model catastrophic outcomes which are triggered by temperature thresholds. A technical point worth mentioning is that while the uncertainty lies in the tipping point, the state varying threshold temperatures in equation (10) are assigned temperature values associated with *time*. This is shown in equation (11) in which $\bar{T}_{AT}(s)$ equals atmospheric temperature value at time $t = s$ on the *never* path, to represent tipping points that *initialize* catastrophes in each SoW. In Figure 1., $\bar{T}_{AT}(s)$ corresponds to atmospheric temperature at nodes $[1, 2, 4]$. In the next section, I introduce the concept of *Stochastic Programming* and further elaborate on why SoWs are assigned to the *time* of catastrophic outcomes rather than tipping points themselves.

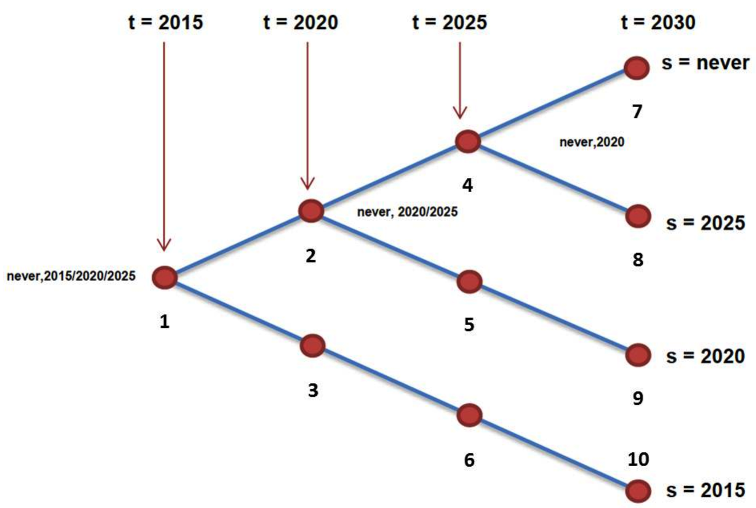

Figure 1. Stochastic Control Recourse Tree

Stochastic Programming with Recourse: Act Then Learn

I have so far introduced the concept of recourse to DICE. In this section, I demonstrate how this recourse tree is used to deal with uncertainty in climate tipping points.

The act-then learn nature of DICESC's stochastic problem motivates the planner to make economic decisions for the future without knowledge of when, or at which tipping point the climate catastrophe will take effect. Let us consider a constant hazard rate $\bar{hr}(t)$, the conditional probability of a catastrophic outcome during period t assuming there has been no such catastrophe to that point in time. Hence for each time period, the planner is able to hedge against the risk of experiencing a catastrophe in the immediate next period given the hazard rate. In the next period the planner realizes one of two states; the planner either faces a state of catastrophe, or finds itself continuing down the never path, yet to encounter the problem of hedging against *the next* period's uncertain outcome (*act then learn*).

Given a constant hazard rate and a recourse tree, the planner assigns a probability to a each SoW path. The conventional hazard rate is mathematically expressed as:

$$hr(t) = \frac{f(t)}{1 - F(t)} \tag{12}$$

where $f(t)$ denotes the probability density of a hazard and $F(t)$, the cumulative distribution of a hazard at time t. As the hazard rate at time t is defined as the probability of a catastrophic outcome conditional on the probability that no catastrophe has taken place before t, we can allocate a probability $\Pi(s)$, to state of the world s by:

$$\Pi(s) = \Gamma(t) \cdot hr(t), \quad t = s \tag{13}$$

where $\Gamma(t)$ is the *survival rate*, the probability of no climate tipping until time $t = s$. By assigning a probability to each deterministic path in which the planner has perfect knowledge of when the catastrophe takes place, the economy is able to hedge against the uncertainty of risky outcomes throughout the entire stochastic horizon.

Again, I come back to the model assumption that uncertainty lies in tipping points. *Then why not assign probabilities directly to tipping points?* By assigning probabilities to tipping points, the model formulates deterministic states of the world in which the temperature threshold is known. Optimizing each deterministic scenario thus becomes a *learn-then act* problem. This problem formulation is easily implemented, but cannot properly model the hedging behavior against risky outcomes, especially when recourse options are available.

Instead, by dealing with uncertainty under a stochastic programming (act-then learn) framework, the planner learns of the tipping point distribution at every time period. Let us go back to the simple recourse tree in figure 1. For instance, given the hazard rate in 2020, the planner hedges against the risky outcome at 2015, only to realize at 2020, the world is still on the no catastrophe SoW. The planner consequently learns that the tipping point is not located at $T_{AT}(2015)$ and lives further to hedge against next period's risky outcome. In contrast, if the planner were to realize in 2020 that the world is at node [3], it learns that $T_{AT}(2015)$ is the threshold temperature that triggers the catastrophic outcome. The stochastic programming framework allows for a gradual resolution of uncertainty in climate temperature thresholds, $\bar{T}_{AT}(s)$.

In *Figures 2.1 and 2.2*, I use DICESC results to represent decision analyses resulting from a stochastic programming formulation with recourse. Note that each branch of the tree denotes a state of the world; e.g. 2030 represents the SoW in which climate tipping occurs at year 2030 and is realized at year 2035. The act-then learn framework is presented in the recourse structure; all states of the world share the *never* branch path until the true climate state is realized.

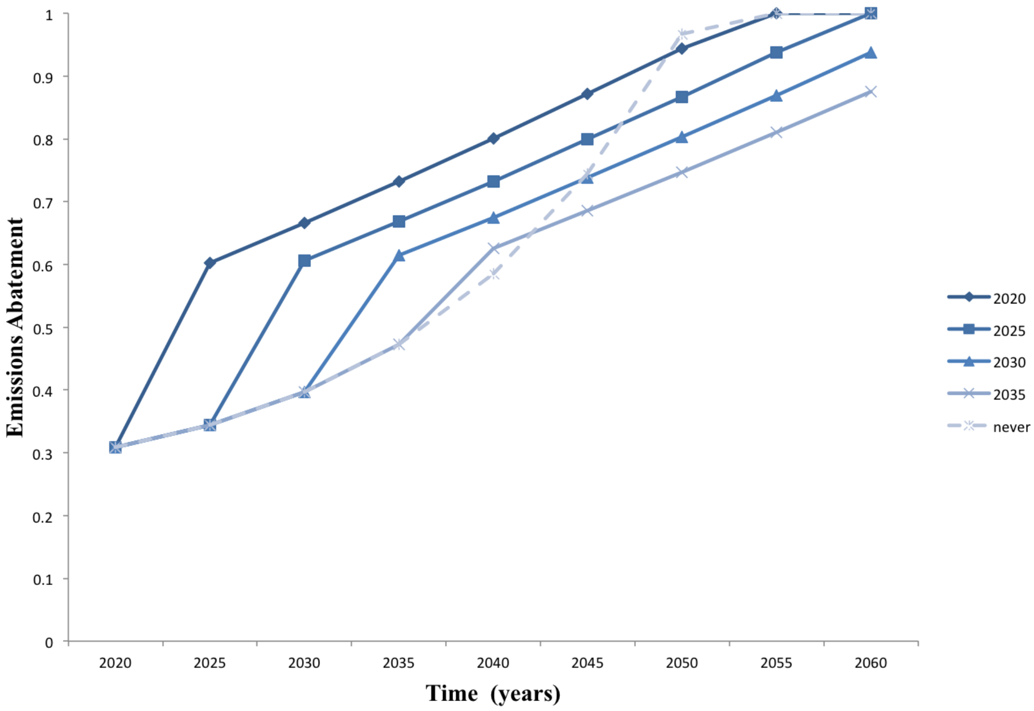

Figure 2.1. Optimal Near-term Abatement in an Act-then Learn Framework

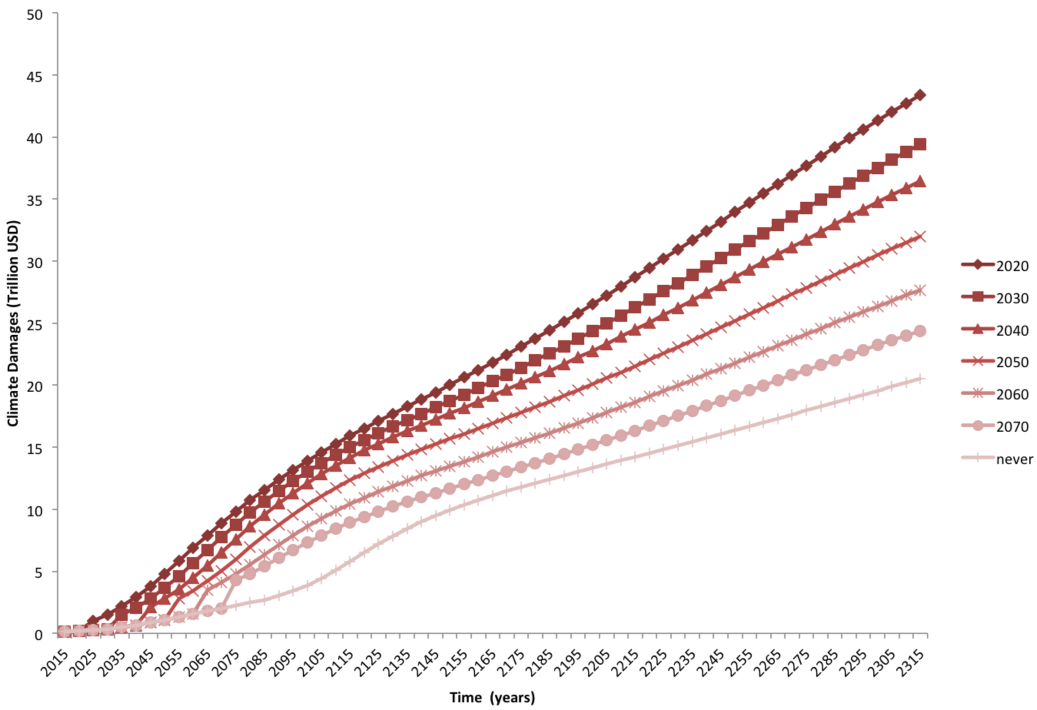

Figure 2.2. Optimized Damages (trillion USD) in an Act-then Learn Framework

Stochastic Control: Endogenous Learning & Controlled Hazard Rates

In the previous framework of stochastic programming, the planner can be characterized as executing *passive learning*. Although the planner progressively learns of the tipping point distribution at each time point in the *never* path, stochastic programming does not allow for the hazard rate to be updated accordingly and learning remains ineffective. Following Lemoine and Traeger 2013, I adopt a simple prior; I now assume that the planner is informed of the distribution domain of the tipping point, $\bar{T}_{AT}$. Note that the maximum temperature threshold in the domain, $\bar{T}_{max}$, is the temperature at which the planner believes with certainty that the climate catastrophe occurs.

$$\bar{T}_{AT} \in [\bar{T}_0, \bar{T}_{max}] \tag{14}$$

The planner's prior belief in the domain of tipping points is stated in (14). Learning down the never path sets the domain in which uncertainty is distributed at time t to be:

$$\bar{T}_{AT} \in [T_{AT}(never, t), \bar{T}_{max}] \tag{15}$$

Why consider temperatures that the mean global atmospheric temperature has already exceeded and are known to have failed to trigger a catastrophic outcome? This active learning is thus incorporated into the hazard rate:

$$h(t) = \frac{F(\bar{T}_{AT} > T_{AT}(t+1)) - F(\bar{T}_{AT} > T_{AT}(t))}{F(\bar{T}_{AT} > \bar{T}_{max}) - F(\bar{T}_{AT} > T_{AT}(t))}. \tag{16}$$

When assuming a uniform prior as does Lemoine and Traeger (2013), the (16) is equivalent to:

$$h(t) = \frac{T_{AT}(t+1) - T_{AT}(t)}{\bar{T}_{max} - T_{AT}(t)}. \tag{17}$$

The hazard rate expression in expression (17) provides an intuitive explanation to the role of learning. Temperature increase incurred between times $[t, t+1]$ heightens the probability of experiencing a catastrophe at the end of the time period. The abatement level at time t thus influences the likelihood of risky outcomes tomorrow and alters the probabilities assigned to each SoW in the recourse tree. While stochastic programming could only justify mitigatory abatement to minimize threshold damages, stochastic control provides reason to abate today as to change the likelihood of the occurrence of risky outcomes.

Although not entirely separate from the numerator's implications of the stochastic control hazard rate, the denominator stresses the Bayesian update aspect of stochastic control. As the global atmospheric temperature reaches higher temperatures without triggering a catastrophe, the prior belief of the tipping point distribution is updated to gradually resolve the uncertainty in climate tipping points (given that $T_{AT}(t+1) > T_{AT}(t)$). Lastly, note that a smaller tipping point domain (as temperatures increase along the never path) bears a greater the likelihood of a catastrophic outcome in response to a unit increase in atmospheric temperature.

Stochastic Control vs. Stochastic Programming

High-level problem formulation for stochastic control and stochastic programming can be conveniently written as a function of abatement.

Stochastic Programming

$$\max_x \sum_s \Pi_s \cdot F_s(x) \tag{18}$$

Equation (18) characterizes the stochastic programming formulations in which Π_s denotes the probability assigned to state of the world s and x, the vector of abatement trajectories for the entire time horizon. $F_s(x)$ then can be viewed as the Lagrangian for each state of the world that encompasses the welfare objective $W(s)$ and the numerous constraints that reflect the social cost of abatement. Note that abatement level x only enters the problem through state indexed Lagrangians. The optimal abatement vector x satisfies the following first order condition:

$$\underbrace{\sum_s \Pi_s \frac{dF(x)}{dx}}_{\text{damage minimizing incentive}} = 0 \tag{19}$$

Hence the optimal policy in stochastic programming is one that optimizes expected welfare, given a fixed set of probabilities for each state of the world. In other words, the optimal abatement strategy is solely characterized by damage minimizing efforts in response to a predestined likelihood of risky outcomes.

Stochastic Control

$$\max_x \sum_s \Pi_s(x) \cdot F_s(x) \tag{20}$$

In stochastic control however, abatement x enters both the state likelihood and the state welfare. In a stochastic control formulation, optimal abatement satisfies:

$$\underbrace{\sum_s \Pi_s(x) \frac{\partial F_s(x)}{\partial x}}_{\text{damage minimizing incentive}} + \underbrace{\sum_s F_s(x) \frac{\partial \Pi_s(x)}{\partial x}}_{\text{precautionary incentive}} = 0 \tag{21}$$

Equation (21) shows that optimal abatement in a stochastic control setup is hence motivated by two incentives; a precautionary incentive to delay or even avoid catastrophes by controlling the hazard rate; and an incentive to minimize catastrophic damages given the probabilistic risky outcomes. Policy optimizes the likelihood of states such that damage minimization results in the highest possible welfare.

3. Results

This section shows results obtained from using DICESC to model the *Arctic ice-sheet melting (ASI)*. The Arctic sea-ice melting (ASI) is characterized to be set off by an increase in global atmospheric temperature between $[0.5, 2]$ °C with a rapid 10 year transition timescale to a new stable equilibrium; among the policy-relevant tipping elements discussed in Lenton, Held, Kriegler, Hall, Lucht, Rahmstorf, and Schellnhuber (2008), Lenton et al. show that the majority of climate tipping such as the Greenland ice sheet melting or the Atlantic thermohaline circulation (THC), are characterized with a significantly longer transition time of over 100 years. The irreversible catastrophic outcomes of ASI are known mainly to be of amplified warming; the absence of ice exposes a darker ocean that absorbs more radiation (Lenton, Held, Kriegler, Hall, Lucht, Rahmstorf, and Schellnhuber, 2008). Lenton et al. further conclude that the strong non-linearity in the decrease of sea-ice coverage may be indication that tipping has already taken place. Considering the 5 year time period used in DICE, and the assumption in DICESC that climate catastrophes are realized in the period after they take place, the Arctic summer sea-ice melting's relatively short transition timescale rendered Arctic sea-ice melting a good fit for the model.

The stochastic horizon is fixed to be 80 years, making the last SoW to take place at year 2095. Additionally, I let $\bar{T}_{AT}$ denote the increase in global atmospheric temperature since year 1900. Lastly, we set the distribution domain for the temperature increase threshold to be $[0.8, 2.0]$ °C. Note that $\bar{T}_0$ corresponds to the initial increase in atmospheric temperature since 1900, which is parameterized as 0.83°C in DICE.

The running time for DICESC with $\bar{T}_{max} = 2.0$ with a stochastic horizon of 80 years is under 3 minutes using NLP GAMS solver *CONOPT* on a laptop computer.

Optimal Abatement

This section compares three optimal abatement policies.

- **DICE2013**
- **DICESC**: DICE stochastic control with $\bar{T}_{max} = 2$
- **DICESP**: DICE stochastic programming

DICE2013 consists of the original version of DICE that does not assume catastrophic damages. For DICESC, I assume a uniform distribution for the tipping points with a distribution domain of $[0.83, 2.0]$. To perform comparative analysis between DICESC and DICESP however, I equate DICESP state probabilities to the optimal probability assignments in DICESC; i.e. $\Pi_s = \Pi_s^*(x) \quad \forall s$. This results in:

$$\underbrace{\sum_s \Pi_s \frac{\partial F_s(x)}{\partial x}}_{\text{DICESP}} = \underbrace{\sum_s \Pi_s^*(x) \frac{\partial F_s(x)}{\partial x}}_{\text{DICESC}} \tag{22}$$

which allows for the comparison of the two runs to assess the additional abatement promoted by precautionary abatement incentives; abatement that arises only when the hazard rate is endogenous. Lemoine and Traeger (2013) show the decomposition of optimal abatement using first order derivative equations similar in nature to equations $(19) - (21)$, and denote the contribution of endogenous hazard rates, the *marginal hazard effect*. I demonstrate this decomposition using numerical solutions in *Figure 3.*.

The difference in abatement levels between DICESP and DICE2013 alone displays the increase in near-term abatement when uncertainty of catastrophic risks are taken into account. The discrepancy represents the planner's hedging strategy to minimize impacts of the risky outcome. *Figure 3.* demonstrates this point; stochastic programming shows a steady increase in the expected abatement path to control for future catastrophic climate damages but shows no sign of precautionary abatement, since it cannot help delay or avoid catastrophic outcomes. DICESC however demonstrates that in a world in which hazard rates can be influenced through temperature control, policy commits significantly more to near-term abatement. In *figure 3.* I display the portion of optimal abatement specifically used to alter hazard rates; i.e. abatement motivated by precautionary incentives.

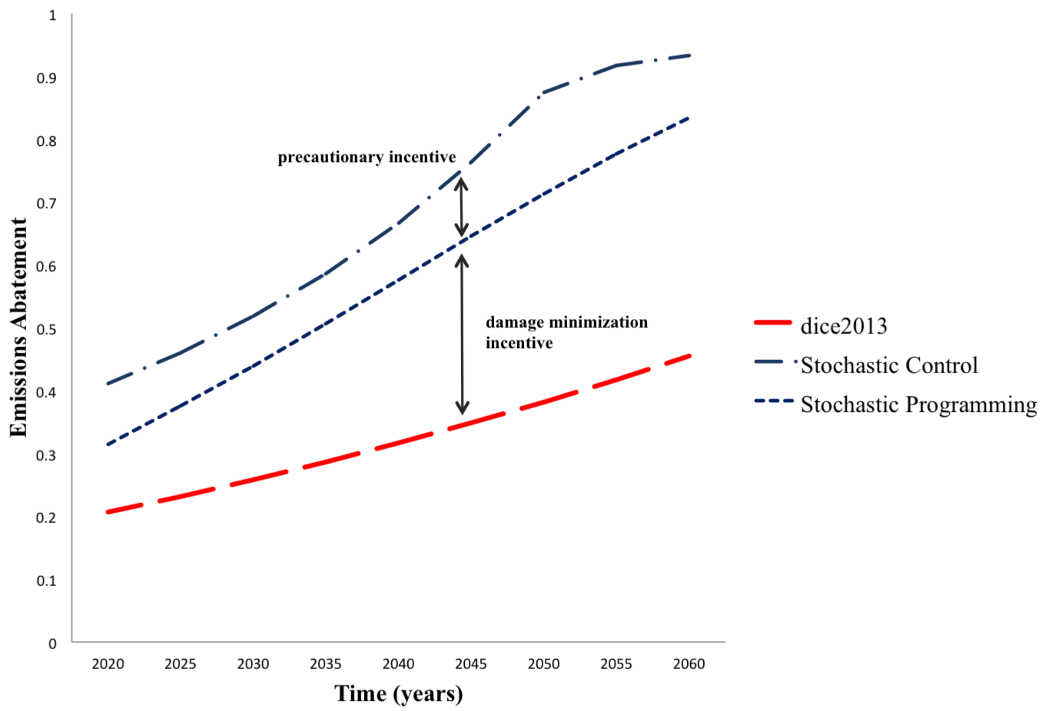

Figure 3. Optimal Near-Term Abatement

4. Sensitivity Analysis

So far I have demonstrated the increase in optimal abatement due to precautionary abatement motives under the risk of a tipping catastrophe using DICESC. *But how robust are the results?* In this section I perform sensitivity analyses to the optimal control trajectory and demonstrate its strong dependence on the distribution of uncertain climate thresholds. The analysis reveals that both damage minimizing and precautionary abatement incentives are susceptible, further stressing the knowledge value of the prior distribution. I conclude the section by assessing the value of perfect information of the tipping threshold for a range of distributions assumptions.

Does Uncertainty Increase Abatement?

I note that the notion of *robustness* is difficult to characterize for decision analyses with century long time horizons. Answering the question at hand serves as a quick but informative measure for policy robustness. Using DICESC, I follow Keller et al.'s approach of comparing the optimal policy with parameter uncertainty to the policy without parameter uncertainty. The conclusions are then compared with that of Keller, Bolker, and Bradford (2004) to demonstrate that the limited efficacy of precautionary abatement in a classical act-then learn framework leads to significantly undervaluing the role of parameter uncertainty to optimal policy.

Keller et al. compares the expected optimal abatement path with a deterministic abatement path in which the tipping point is believed with certainty to be at the distribution's *expected tipping point*. Assuming a uniform distribution over the tipping point domain $[0.8, 2]$ °C, the expected tipping point becomes 1.4 °C. I respectively determine the state of the world s^* that maximizes social welfare W assuming climate tipping at the expected tipping point $E[\bar{T}_{AT}]$ and conclude $s^* = 2045$ for $\bar{T}_{AT} = E[\bar{T}_{AT}] = 1.4$. *Figure 4.* below compares the optimal stochastic control abatement path assuming a uniform distribution to the deterministic optimal abatement path assuming $\bar{T}_{AT} = 1.4$. The DICESC optimal abatement path given $\bar{T}_{max} = 2.0$ solves the problem:

$$\max_x \sum_s \Pi_s(x) \cdot F_s^{2.0}(x), \tag{23}$$

and is compared to the optimal policy neglecting uncertainty, which in turn solves the following deterministic problem:

$$\max_x F_{s^*}^{E[\bar{T}_{AT}]}(x). \tag{24}$$

In this comparative setting, (24) corresponds to:

$$\max_x F_{2045}^{1.4}(x). \tag{25}$$

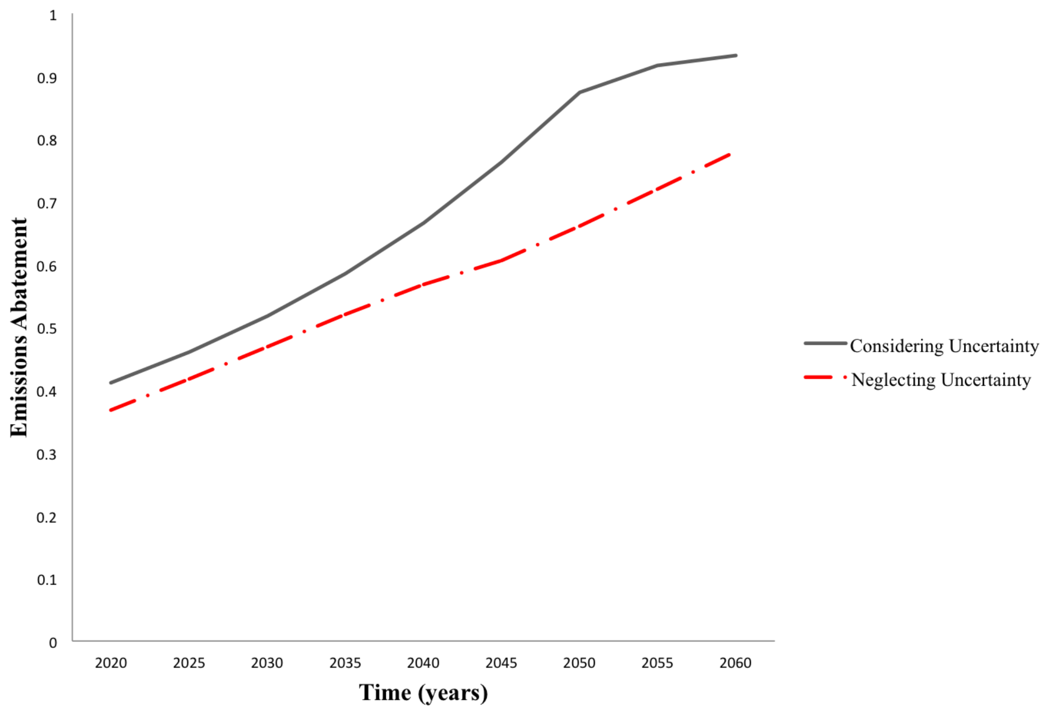

Figure 4.1. Decision Analysis Under Tipping Point Uncertainty

While Keller et al. concludes that parameter uncertainty decreases optimal abatement, DICESC suggests otherwise; under a uniform tipping point distribution on the domain $[0.8, 2.0]$ °C and a deterministic threshold damage function, I show that uncertainty can *increase* optimal abatement. Since Keller, Bolker, and Bradford (2004) further incorporates uncertainty in threshold damages, DICESC also implements stochastic optimization assuming low threshold damages; in this run, the order of the threshold damage term is changed from cubic to linear. The previous comparative analysis is applied to the case of uncertain tipping points and low threshold damages, only to prove moderate robustness of the conclusion previously obtained (*Figure 4.2.*).

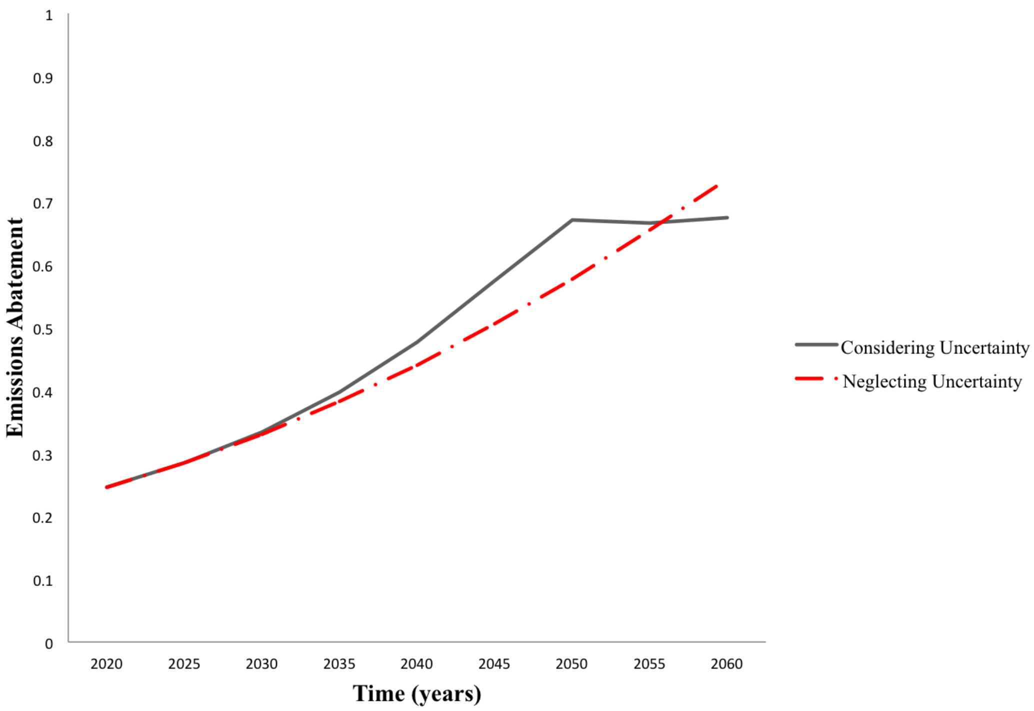

Figure 4.2. Decision Analysis Under Tipping Point Uncertainty & Low Threshold Damages

So was Keller, Bolker, and Bradford (2004) wrong to conclude that uncertainty decreases optimal abatement? Keller et al. base their conclusions on specific distributive assumptions about the uncertain parameters; climate sensitivity is assumed to be normally distributed between 1.5 and 4.5 degrees and threshold damages, uniformly distributed between 0 and 3 percent of Gross World Product. Also note the model incorporates no scope for endogenous hazard rates. Based on these presumptions Keller et al.'s conclusions remain valid. However in a formulation in which optimal abatement is susceptible to the assumed distribution that characterizes uncertainty, Keller's conclusion only represents a fragmentary assessment of the question at hand.

To demonstrate the constrained validity of their claim, I replicate Keller et al.'s conclusions using a stochastic programming formulation of the problem. *Figure 5.1.* shows that the deterministic optimization path dominates the optimal abatement trajectories obtained using a stochastic programming framework, regardless of the assumed distribution priors. Less abatement in the stochastic run compared to the deterministic intuitively implies that optimal abatement is high for the expected tipping point, but decreases significantly as we move toward the right tail of the distribution (Keller, Bolker, and Bradford, 2004). By assuming a uniform tipping point distribution then, the social planner's decisions recognize that all tipping points in the domain are equally likely to trigger a catastrophe. Without control over hazard rates in a deterministic setting, the welfare benefits from moving right of the expected tipping point outweigh the costs of moving left, resulting in an act-then learn stochastic abatement trajectory that decreases with parameter uncertainty. Further deterministic runs however show that this dominance is limited to maximum tipping points ($\bar{T}_{max}$) within the domain $[2, 4]$ °C (*Figure 5.2.*).

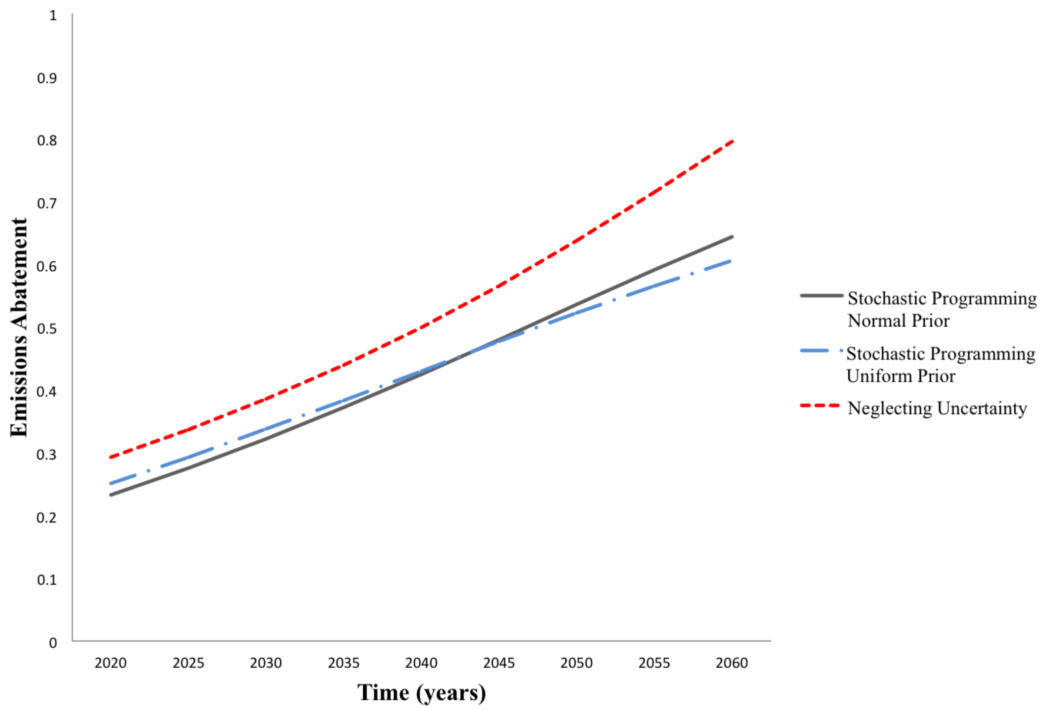

Figure 5.1. Replicating Keller et al.'s Conclusions in a Stochastic Programming Framework

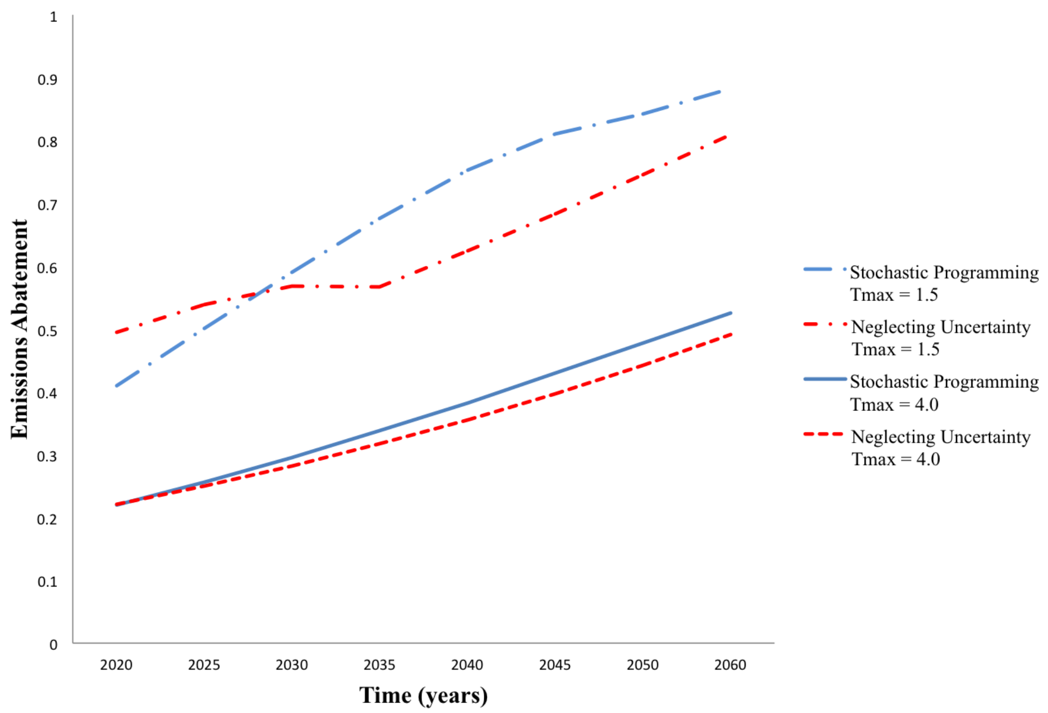

Figure 5.2. Robustness Check for Keller et al.'s Conclusions

Cases for $\bar{T}_{max} \geq 4.0$ or $\bar{T}_{max} \leq 2.0$ generate contrasting results. When assuming $\bar{T}_{max} \geq 4.0$, the expected tipping point takes a value greater than 2.4 °C, for which the economy's cost-benefit analysis of abatement chooses to completely avoid climate tipping. This decision is prompted by the fact that controlling atmospheric temperature increase below 2.4 °C does not impose stringent constraints on economic production. Considering tipping point uncertainty for $\bar{T}_{max} \geq 4.0$ consequently results in an increase in abatement (*Figure 5.2.*). The case of $\bar{T}_{max} \leq 2.0$ gives rise to the same outcome for different reasons. The decision analysis for $\bar{T}_{max} = 1.5$, for which the expected tipping point takes value approximately 1.2 is displayed in *Figure 5.2.*. In this case a scenario in which the economy lets go of temperature control and allows tipping to take place in year 2035 is optimal, suggesting that the cost of controlling temperature increase below 1.2 °C is unsustainable.

Although strictly speaking the results of DICESC and FRANC are not directly comparable, a simple stochastic programming formulation that replicates Keller et al.'s conclusion, reveals the lack of robustness of *damage minimizing* optimal abatement to the distribution domain of tipping thresholds. Nevertheless, the optimal abatement decomposition in *Figure 3.* shows that the abatement deficiency in FRANC is largely attributed to the absence of *precautionary* abatement, motivated by the endogeneity of the hazard rate. The next section further analyzes the sensitivity of *precautionary* abatement to stress the relevance of distributive assumptions for a robust, comprehensive assessment of optimal abatement.

Decomposing Precautionary Abatement

The decomposition of optimal abatement presents precautionary abatement as abatement in surplus of optimal abatement under stochastic programming. Since this surplus abatement is a direct product of endogenous hazard rates, sensitivity of precautionary abatement can be presented by decomposing the effects of abatement to hazard rates.

Following Lemoine and Traeger (2013), DICESC defines the endogenous hazard rate as a function of atmospheric temperature and the cumulative distribution function of tipping points. I take the total derivative of the hazard rate with respect to abatement to obtain:

$$\frac{dhr(F(T_{AT}))}{dx} = \underbrace{\left(\frac{\partial hr(F)}{\partial F}\right)\left(\frac{\partial F(T_{AT})}{\partial T_{AT}}\right)}_{(1)}\underbrace{\left(\frac{\partial T_{AT}(x)}{\partial x}\right)}_{(2)}. \tag{26}$$

While term (2) corresponds to the sensitivity of atmospheric temperature to abatement, which directly relates to *climate sensitivity*, term (1) represents the *hazard rate sensitivity*; the hazard rate's reponsiveness to changes in atmospheric temperature. Note that while the climate sensitivity inevitably affects the optimal abatement path regardless of whether uncertainty is taken into account, the hazard rate sensitivity exclusively affects precautionary optimal abatement under uncertainty.

I perform sensitivity analyses to term (1), to elucidate the pathways in which the distribution of tipping points influence the level of precautionary abatement. To decrease the value of the maximum temperature threshold means to assume a sure event of climate tipping at a lower temperature and as a result, a unit increment of warming bears a greater risk of hazard. So overall, a lower $\bar{T}_{max}$ increases *hazard rate sensitivity* and gives rise to more precautionary abatement, under the assumption of a constant *climate sensitivity* effect. *Figure 6* displays the optimal precautionary abatement levels at year 2020 for different $\bar{T}_{max}$ values, assuming a truncated normal threshold distribution. It displays decreasing levels of precautionary abatement as the threshold distribution is spread across a greater temperature domain, and demonstrates the sensitivity of abatement to the domain of the threshold distribution.

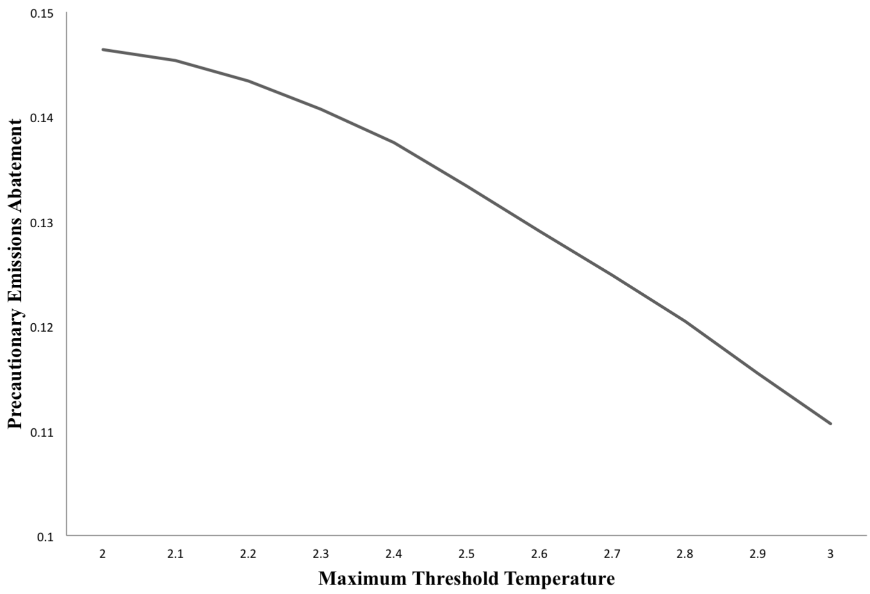

Figure 6. Optimal Precautionary Abatement Levels at Year 2020

Hazard rate sensitivity is also affected by the assumed class of the threshold distribution. To assess the sensitivity of DICESC to the assumed tipping point distribution, the first and second moments of the truncated normal distribution are fixed to be identical to those of the previously applied uniform distribution. *Figure 7.* displays the truncated normal tipping distributions obtained by calibrating moments to the corresponding uniform distributions.

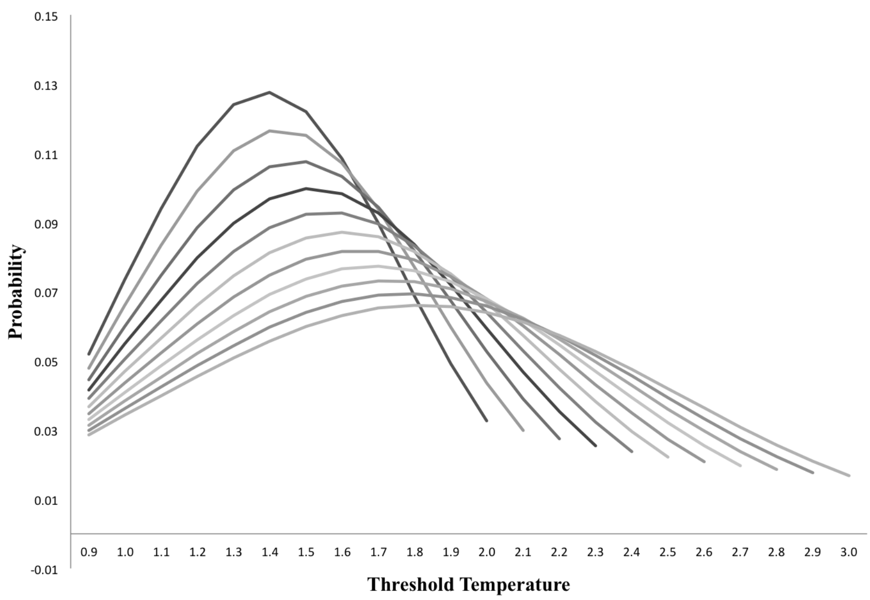

Figure 7. Truncated Normal Threshold Distributions

I display the difference in abatement levels between assuming a truncated normal and uniform tipping distribution, in *Figures 8.1. & 8.2.*. The discrepancy in expected abatement can be explained intuitively by characterizing abatement trajectories at each end of the distribution tail. The act-then learn formulation of the problem requires the planner to consider all deterministic states of the world ranging from a disastrous scenario in which climate tipping occurs at temperatures as low as 0.9 °C, to ones in which climate tipping can be avoided with ease. While the left-end of the tail is often characterized by non-smooth abatement trajectories that represent drastic measures of mitigation, less stringent thresholds located toward the right of the distribution constitute smooth convex or concave policies to exhibit a gradual increase in temperature control.

When assuming a truncated normal distribution, the extreme scenarios associated at each end of the tail are not given much weight, allowing for a steady increase in expected optimal abatement throughout time. The steady decrease in precautionary abatement is caused by the gradual convergence between stochastic programming and stochastic control optimal control paths, which represents the increasing role of *damage minimizing* abatement relative to precautionary abatement.

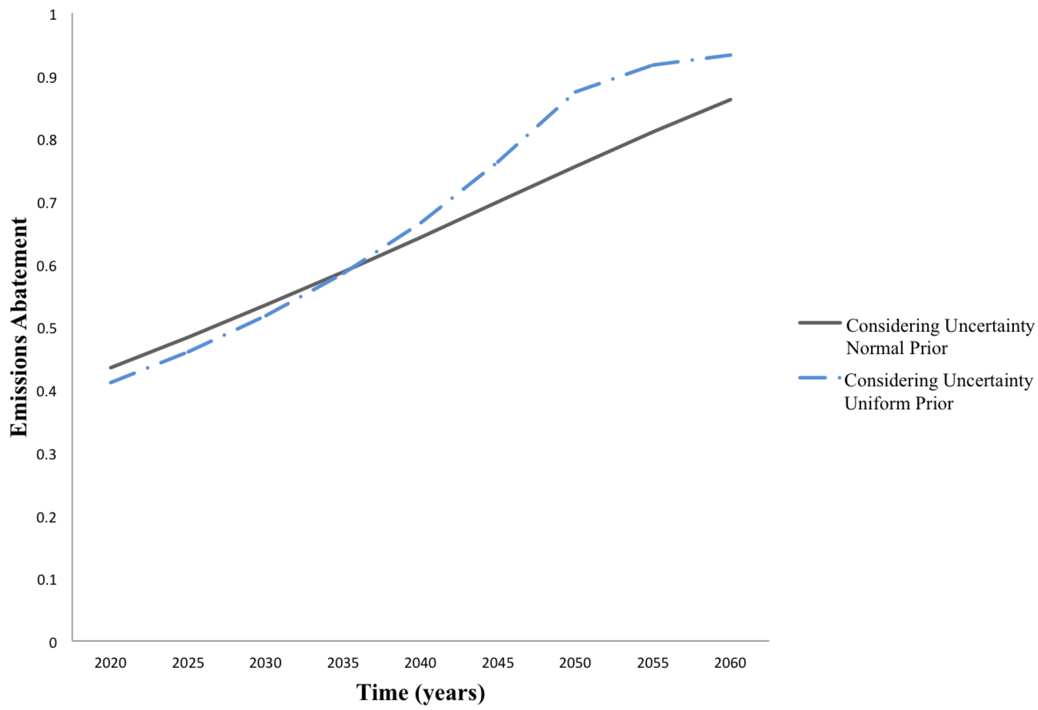

Figure 8.1. Optimal Abatement at $\bar{T}_{max} = 2\,°\mathrm{C}$ for Two Threshold Distributions

Assuming a uniform distribution, equal weight is imposed on all tipping point scenarios including the tails of the distribution, which is reflected at each end of the stochastic trajectory. On one hand, significant weight put on the right end of the distribution gives rise to less optimal abatement in the immediate couple of decades; analysis proves that for thresholds above 1.5 °C, immediate near-term abatement decreases at an increasing rate. Moreover, for certain disastrous states that assume tipping at $\bar{T}_{AT} \leq 1.5\,°\mathrm{C}$, cost-benefit analysis of abatement proves that it is economically beneficial to allow tipping catastrophes to occur as early as the year 2030, since temperature control below the threshold becomes economically unsustainable. In such scenarios, precautionary abatement is kept at moderately low levels in the early stages.

On the other hand, significant weight allocated to the left-end of the tail generates high abatement levels later in the time horizon. Although this result is mainly achieved through stringent temperature control in the lower-end of the distribution, a significant portion of precautionary abatement is attributed to scenarios associated with less stringent thresholds;

further analysis shows that efforts to avoid threshold damages with a greater likelihood generates convex precautionary abatement paths even for thresholds located at the right-end of the tail. Again, the decrease in precautionary abatement is explained by the convex-*damage minimizing* abatement paths that lessens the output difference between stochastic control and stochastic programming.

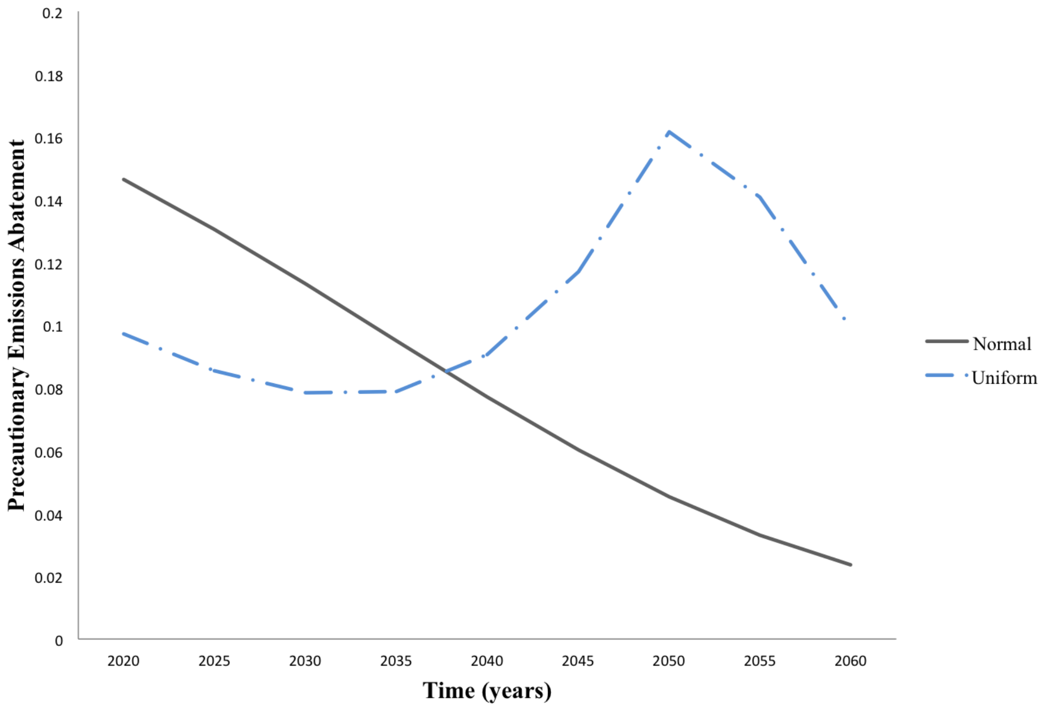

Figure 8.2. Optimal Precautionary Abatement at $\bar{T}_{max} = 2\,^{\circ}\mathrm{C}$ *for Two Threshold Distributions*

Assessing the Expected Value of Perfect Information (EVPI)

The discussion thus far brings to light key factors that determine optimal abatement under tipping point uncertainty; namely, the responsiveness of hazard rates to the economy's mitigatory efforts and the assumed distribution of climate thresholds. The distinct pathways in which threshold distributions influence components of optimal abatement reveals the optimal abatement policy's sensitivity to distributive assumptions under the risk and uncertainty of climate tipping. Lastly, I assess the *Expected Value of Perfect Information (EVPI)* in DICESC; the value of instantaneously resolving all uncertainty to obtaining perfect information of the temperature threshold. The results in this section illustrate the drastic decrease in the EVPI through means of learning and endogenous hazard rates, for which the levels again marginally vary depending on the distribution of tipping points. I adopt the approach presented in Manne and Richels (1992) in computing the EVPI index and present the results in *Figure 9.*.

The EVPI is computed by comparing the expected welfare resulting from optimization under uncertainty, to the expected welfare under perfect information:

$$EVPI_{\bar{T}_{max}} = \int_{\bar{T}_0}^{\bar{T}_{max}} W^{*}_{\bar{T}_{AT}} f(\bar{T}_{AT})\, d\bar{T}_{AT} - W^{U}_{\bar{T}_{max}} \tag{27}$$

$$\begin{aligned} \text{where} \quad W^*_{\bar{T}_{AT}} &= W(s^*) \\ W^U_{\bar{T}_{max}} &= \sum_s \Pi_s(x^*) \cdot F_s(x^*) \\ \text{for} \quad x^* &= \arg\max_x \sum_s \Pi_s(x) \cdot F_s(x) \end{aligned} \tag{28}$$

for which I compute the discretized formulation:

$$EVPI_{\bar{T}_{max}} = \sum_{\bar{T}_0}^{\bar{T}_{max}} W^*_{\bar{T}_{AT}} f(\bar{T}_{AT}) - W^U_{\bar{T}_{max}} \tag{29}$$

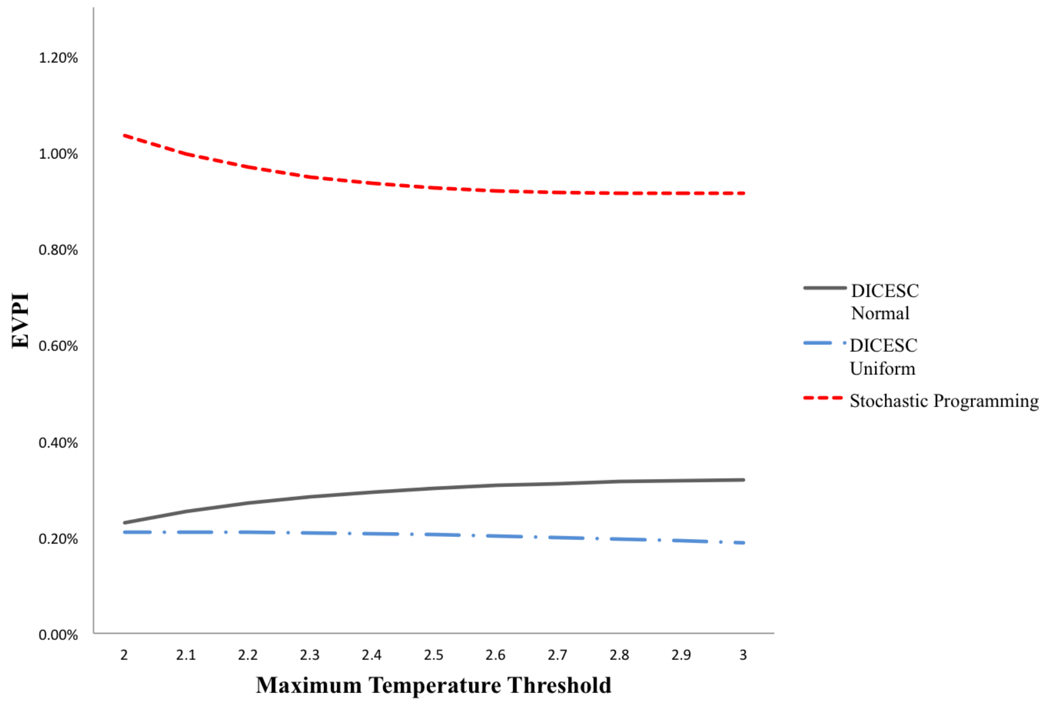

Figure 9. Expected Value of Perfect Information

Note that $f(\cdot)$ is the pdf of the initial prior and that the EVPI percentage index in *Figure 9.* is expressed as a ratio over the expected utility under uncertainty, $W^U_{\bar{T}_{max}}$. The figure illustrates two important points. Firstly, while the EVPI value resulting from fixed hazard rates (stochastic programming) hovers around one percent of expected welfare, the EVPI index generated by DICESC displays values that are far less significant. The discrepancy represents the *value of learning* adopted in the stochastic control setting. In stochastic control, endogenous learning gradually resolves uncertainty and also drives efficient hedging by means of precautionary abatement; hence the gains from obtaining perfect information are not extensive. A stochastic programming formulation however incorporates no scope for learning, and the value of resolving uncertainty is considerable as the means of hedging against the risky outcome is only limited to *damage minimizing abatement*. Because the likelihood of experiencing catastrophic damages are fixed despite the economy's mitigatory actions, expected threshold damages are significantly greater than those resulting from stochastic control; this adds more value to resolving uncertainty.

Although the EVPI indices are reasonably robust to the assumed threshold distributions, there is a notable difference in the trend of the two EVPI plots. First note that as $\bar{T}_{max} \longrightarrow \infty$, the two expected welfare paths will eventually converge. The upward trend when assuming a truncated normal tipping point distribution then represents the slow convergence of expected welfare under uncertainty to that under perfect information. This can be explained intuitively by the properties of marginal increase in expected welfare with respect to $\bar{T}_{max}$ under each distribution. A unit increase in $\bar{T}_{max}$ puts more weight on the right-end of the prior under a uniform distribution relative to a truncated normal distribution. Since higher temperature thresholds are characterized with higher welfare, expected welfare under uncertainty subsequently increases at a faster rate when assuming a uniform threshold distribution. Despite the contrasting rate of convergence between the two expected welfare trajectories under uncertainty, expected welfare under perfect information remains relatively robust to tipping point distribution which brings about the distinct EVPI trends displayed in *Figure 9.*.

5. Conclusion

The research demonstrates the role of precautionary abatement when hedging against tipping catastrophes. Using a stochastic control version of DICE, I show that the possibility of a catastrophic outcome in the future increases abatement levels not only to minimize catastrophic damages, but also to delay or even avoid the occurrence of such outcomes. I decompose optimal abatement with respect to mitigatory incentives and numerically illustrate the precautionary abatement efforts that can *increase* optimal abatement under uncertainty. Sensitivity analyses expose the policy's heavy dependency on the distribution of uncertain climate thresholds. A decomposed analysis of abatement reveals that both damage minimizing and precautionary abatement incentives are susceptible to assumed distributive properties, further emphasizing the knowledge value of the prior distribution for robust policy assessment. Lastly, an EVPI analysis shows that endogenous hazard rates and learning significantly decreases the value of perfect information regarding the uncertain temperature threshold.

Modeling risk and uncertainty is more than merely simulating catastrophic damages and parameter uncertainty; stochastic control points to sensible representation of learning and decision making within a stochastic setting, as yet another critical aspect that models cannot afford to overlook. Although the IAM literature on the representation of catastrophic climate outcomes is too underdeveloped to benchmark a well-rounded model, models that omit such integral elements in modeling climate uncertainty cannot provide a comprehensive decision analysis. Finally, this research demonstrates that decision analyses under climate uncertainty alone cannot outline a robust hedging strategy. Improved understanding of tipping point distributions and refined assessment of threshold damages by climate researchers is deemed crucial for a comprehensive assessment of optimal policy.

References

Ackerman, F., E. A. Stanton, and R. Bueno (2010): "Fat tails, exponents, extreme uncertainty: Simulating catastrophe in DICE," *Ecological Economics*, 69(8), 1657–1665.

——— (2013): "Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy?," *Environmental and Resource Economics*, 56(1), 73–84.

Backus, D. K., B. R. Routledge, and S. E. Zin (2005): "Exotic preferences for macroeconomists," in *NBER Macroeconomics Annual 2004, Volume 19*, pp. 319–414. MIT Press.

Baker, E. (2009): "Optimal policy under uncertainty and learning about climate change: A stochastic dominance approach," *Journal of Public Economic Theory*, 11(5), 721–747.

Cai, Y., K. L. Judd, and T. S. Lontzek (2012): "DSICE: a dynamic stochastic integrated model of climate and economy," *Center for Robust Decision-Making on Climate Policy, Working Paper*, (12-06).

Defourny, B., D. Ernst, and L. Wehenkel (2012): "Multistage stochastic programming: A scenario tree based approach to planning under uncertainty," *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions*, pp. 97–143.

Dixit, A. K. (1990): *Optimization in economic theory*. Oxford University Press.

Epstein, L. G., and S. E. Zin (1989): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica: Journal of the Econometric Society*, pp. 937–969.

Hope, C. (2006): "The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern," *Integrated assessment*, 6(1).

Kaufman, N. (2012): "The bias of integrated assessment models that ignore climate catastrophes," *Climatic change*, 110(3-4), 575–595.

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.

Kelly, D. L., and C. D. Kolstad (1999): "Bayesian learning, growth, and pollution," *Journal of Economic Dynamics and Control*, 23(4), 491–518.

Lemoine, D., and C. Traeger (2013): "Watch your step: optimal policy in a tipping climate," *American Economic Journal: Economic Policy, forthcoming*.

Lemoine, D. M., and C. Traeger (2010): "Tipping points and ambiguity in the integrated assessment of climate change," *University of California, Berkeley, Department of Agricultural and Resource Economics working paper*, 1111.

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.

Manne, A., R. Mendelsohn, and R. Richels (1995): "MERGE: A model for evaluating regional and global effects of GHG reduction policies," *Energy policy*, 23(1), 17–34.

Manne, A. S., and R. G. Richels (1992): *Buying greenhouse insurance: the economic costs of carbon dioxide emission limits*. MIT Press.

Metz, B., O. R. Davidson, P. R. Bosch, R. Dave, and L. A. Meyer (2007): "Contribution of Working Group III to the fourth assessment report of the Intergovernmental Panel on Climate Change," .

Nordhaus, W., and P. Sztorc (2013): "DICE 2013R: Introduction and User's Manual," .

Nordhaus, W. D. (1994): *Managing the global commons: the economics of climate change*. MIT press Cambridge, MA.

——— (2011): "Integrated economic and climate modeling," Discussion paper, Cowles Foundation for Research in Economics, Yale University.

Pindyck, R. S. (2013): "Climate Change Policy: What do the models tell us?," Discussion paper, National Bureau of Economic Research.

Rutherford, T. F. (2001): "Calibration of Models with Multi-Year Periods," Discussion paper, mimeo.

Stocker, T. F., D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, P. M. Midgley, et al. (2013): "Climate Change 2013. The Physical Science Basis. Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change-Abstract for decision-makers," Discussion paper, Groupe d'experts intergouvernemental sur l'evolution du climat/Intergovernmental Panel on Climate Change-IPCC, C/O World Meteorological Organization, 7bis Avenue de la Paix, CP 2300 CH-1211 Geneva 2 (Switzerland).

Urban, N. M., P. B. Holden, N. R. Edwards, R. L. Sriver, and K. Keller (2014): "Historical and future learning about climate sensitivity," *Geophysical Research Letters*, 41(7), 2543–2552.

Webster, M., N. Santen, and P. Parpas (2012): "An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty," *Computational Management Science*, 9(3), 339–362.

Webster, M. D. (2000): "The curious role of" learning" in climate policy: should we wait for more data?," .