Mathematical Programming Tools to Deal with Uncertainty in Climate Economics

By

Wonjun Chang

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Agricultural and Applied Economics)

at the

UNIVERSITY OF WISCONSIN-MADISON

2017

Date of final oral examination: 4/17/2017

The dissertation is approved by the following members of the Final Oral Committee:

Thomas F. Rutherford, Professor, Agricultural and Applied Economics
Corbett Grainger, Associate Professor, Agricultural and Applied Economics
Daniel Phaneuf, Professor, Agricultural and Applied Economics
Michael Ferris, Professor, Computer Science
Derek Lemoine, Assistant Professor, Economics

To Jiya

Acknowledgements

First and foremost I would like to thank God. None of this would have been possible without his partnership. I express my sincere gratitude to my two advisors Tom Rutherford and Corbett Grainger for their training, guidance and support throughout my graduate school career. Many thanks also to my committee members, Dan Phaneuf, Michael Ferris and Derek Lemoine who kindly offered direction and encouragement in my dissertation writing. I am deeply indebted to my good friend Barbara Forrest, who helped me embark on this journey at UW and cheered me on every step of the way in finishing my degree.

I wish to acknowledge my family for their endless love and support. I thank my loving wife, Sam, and daughter, Jiya, for fueling my motivation whenever I would run out. I have become a stronger and better person because of you. To my brother Drew; I'm blessed to have your company as a brother, both in academia and at home. I am also grateful to my generous and kind-hearted landlords (who are also my parents-in-law), Mark and Becky, for putting up with my family's mess at home. Last but not least, I thank my parents Kitae and Haekyung for their unconditional love. Thank you for always having my back. Your support has been the greatest gift anyone has ever given me.

Abstract

How does risk and uncertainty in climate thresholds impact optimal short-run mitigation? Determining the impact of uncertain climate outcomes on near-term climate policy has long been a subject of debate in the integrated assessment community. In this dissertation I present mathematical programming tools to assess this so called "uncertainty effect" on climate policy.

I first present a minimal *act-then learn* stochastic control model (DICESC) based on Bayesian learning. This captures the sequential process of decision making under uncertainty based on new observational evidence. In an act-then learn setting the possibility of climate tipping in the future increases optimal abatement to delay or avoid threshold damages. While uncertainty increases the incentive for precautionary abatement, increments in mitigation and the expected value of perfect information (EVPI) may not be robust, as these can depend on higher moments of the Bayesian prior. Hazard rate distributions sharing the same mean and standard deviation may have different values of information and optimal short term mitigation policies. Humility is called for as it seems unlikely that we can determine the distribution of risk with sufficient precision.

Insufficient precision in the Bayesian prior necessitates focus on the notion of ambiguity. Given multiple distributions of risk, I implement four models of ambiguity aversion and derive optimal policies for each ambiguity attitude. Results indicate a robust precautionary incentive to increase abatement under ambiguity aversion. Since effective priors are defined as a mixture of the prior set in each period, abrupt transitions between priors can impose

threshold effects on policy even before tipping occurs; an ambiguity averse agent exhibits incentives to delay reaching temperature points at which the assumed Bayesian prior takes a turn for the worse.

In the last chapter, I present a mixed complementarity problem (MCP) formulation of dynamic programming (DP) problems. The MCP approach replaces conventional value function iteration by the solution of a one-shot square system of equations and inequalities. Three numerical examples illustrate my approach and demonstrate that the DP-MCP algorithm can compute equilibria much faster than traditional value iteration. In addition, the MCP approach accommodates corner solutions in the optimal policy.

Table of Contents

Chapter 1.

Introduction

In addressing the economic impacts of climate change, numerical Integrated Assessment Models (IAMs) have assumed an important role by coupling climate models with models of economic growth. The main objective of such climate-economy IAMs is to optimize the trade off between emissions abatement today and climate damages in an uncertain future, to obtain an optimal policy actionable in the near-term. However, given the uncertainties of climate impacts under global warming (let alone the myriad uncertainties of the climate system) the integrated assessment community has in general struggled in determining *uncertainty effects*—how the risk and uncertainty of climate impacts influence the near-term climate policy.

One aspect of climate impacts that sees little room for agreement is the risk and uncertainty of climate catastrophes. Recent literature points out that the most important driver of the social cost of carbon is the likelihood of catastrophic climate outcomes. Neglecting such climate risks can output emissions abatement policies that are seriously misleading [Pindyck, 2013, Weitzman, 2011].

As a result, climate catastrophes triggered when atmospheric temperatures exceed a climate temperature threshold, also known as tipping elements, have entered numerous IAM models. Policy relevant tipping elements include the Atlantic thermohaline circulation (THC) collapse, Boreal forest dieback and Greenland ice sheet (GIS) and Arctic summer sea-ice melting. Incorporating climate tipping to optimal growth models however has proven to be tech-

nically challenging due to the abruptness and non-convex nature (hysteresis effect) of the process. 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 et al., 2004].

Stochastic frameworks used to deal with the risks and uncertainties of climate tipping consist largely of two types; Monte Carlo simulation based and dynamic programming based frameworks. Efforts to realistically represent the complex structure of climate uncertainties and the non-convex character of damages have confined uncertainty analyses to parameterized decision rules, such as Monte Carlo simulations. However, as the literature points out, the Monte-Carlo averaging approach is misrepresentative of how we make decisions under uncertainty. The reason for this is that the approach is, in essence, an averaged sensitivity analyses, with each simulation equivalent to running a deterministic model [Crost and Traeger, 2013]. To devise a policy from with perfect information on the occurrence and impact of tipping essentially takes away the most damaging aspect of uncertain climate outcomes—its *abruptness*. Averaging over these deterministic runs can lead to determining the wrong sign, if not the wrong magnitude of the uncertainty effect.

On the other hand, efforts to create a fully stochastic formulation of abrupt, irreversible climate outcome have steered the IAM community to use dynamic programming (DP) frameworks. Unlike deterministic models, stochastic DP frameworks allow for the model to treat tipping points as intrinsic random events that remain uncertain until catastrophic events are realized. Despite methodological advances in employing dynamic programming to climate models however, conventional dynamic programming is not the most accessible of methods to obtain numerical solutions with. Implementation is indeed computationally burdensome and dimensionally restricted, especially for assessing sequential decision policies characterized by

a large set of actions [Powell, 2007]. In the main chapters of this disseration, I introduce a third type of framework—one based on stochastic programming—to deal with uncertainties including, but not limited to climate tipping.

In chapter 2, I introduce a minimal stochastic framework of precautionary mitigation using a *stochastic control* version of Nordhaus' DICE model [Nordhaus and Sztorc, 2013], *DICESC*. DICESC uses stochastic programming with endogenous state probabilities. It maintains a recourse structure that avoids the scaling limitations of classical multistage stochastic programming, while treating tipping points as an intrinsic uncertainty. Moreover, its formulation as a one-shot optimization problem allows DICESC to be solved using off-the-shelf commercial solvers, replicating insights from recursive methods presented in Lemoine and Traeger [2014] and Cai et al. [2012], in a matter of minutes.

In chapter 3, I discuss the critical importance of *learning* in the decision making process under uncertainty. As previously discussed, Monte Carlo simulations have widely been used to assess the effect of uncertainty to policy, especially in complex integrated assessment models. To handle uncertainty properly (intrinsically), some studies have used *Monte Carlo Fitting*, a strategy that identifies a single policy to optimize over the range of possible outcomes. Although this approach represents decision making under uncertainty, the stochastic framework allows for no scope for learning. In this chapter, I demonstrate the critical role learning plays in determining the optimal policy by comparing DICESC results with that of *lack-of-learning* counterfactuals. I show that the absence of learning, either in the history of the event process, or of the Bayesian prior used to characterize the likelihood of an event, can lead to optimal policies that are severely misleading.

Chapters 2 and 3 deal with the notion of *risk*. The term risk is used to characterize environments in which the probability distribution of a random event taking place is known

with certainty. In contrast, the notion of *uncertainty* (also known as Knightian uncertainty or ambiguity) characterizes an instance in which the decision maker is unsure of the probability distribution that describes the likelihood of the risky event occurring. In chapter 4, I deal with the uncertainty of climate thresholds. Given multiple priors that characterize the risk of tipping, I present how attitudes toward tipping ambiguity can be modeled in DICESC, based on four different definitions of ambiguity aversion. The contributions of this chapter are two fold. First, I discuss *how* ambiguity in the likelihood of climate tipping can be modeled based on four different definitions of ambiguity aversion in a *general* stochastic control setting with *multiple prior distributions*. Secondly, I provide numerical assessments of optimal policy under each ambiguity averse attitude, to add to the optimal policy and uncertainty effect literatures.

Chapter 5 introduces a new framework. In this chapter, I present a mixed complementarity problem (MCP) formulation of DP problems with continuous state space. As mentioned, a sizable portion of the climate-economy modeling literature relies heavily on dynamic programming frameworks when assessing the effects of uncertainty. But unlike the types of optimization problems in mathematical programming, there is no standard formulation of dynamic programming (DP), nor is there an off-the-shelf solver package designated specifically for DP [Brandimarte, 2014]. DP is a principle, and there exist multiple formulations and customizations that focus on solution accuracy and computational efficiency when it comes to implementation [Tauchen, 1986, Cai and Judd, 2015, Judd et al., 2014].

One of the most common computational approaches of formulating DP problems using climate IAMs is *value function iteration.* This method is widely used due to its monotonic convergence properties and straightforward implementation. Despite the algorithm's stability however, the iterative aspect of the diagonalization procedure makes implementation computationally burdensome and time consuming. This problem quickly worsens in grid size

and state dimensions, often prompting the modeler to devise technically creative ways to get around the problem of intractability.

The objective of using the MCP approach is to replace conventional value function iteration by the solution of a one-shot square system of equations and inequalities. The MCP approach is designed to significantly reduce the run-time required to solve the DP problem, and more importantly, allow for the proper treatment of corner solutions while solving for the optimal policy in dynamic programming. I include three numerical examples to illustrate the approach and demonstrate the advantages of using *DP-MCP*, including a multisector stochastic growth model and a hydro-power scheduling problem.

Chapter 2.

Catastrophic Thresholds, Bayesian Learning and the Robustness of Climate Policy Recommendations

1. Introduction

In addressing the economic impacts of climate change, numerical Integrated Assessment Models (IAMs) of varying complexity have assumed a central role in summarizing climate science (i.e., the anticipated impact of greenhouse gases on atmospheric temperature) and economic damage. IAMs coupling economy and climate include DICE [Nordhaus and Sztorc, 2013, Nordhaus, 1994], FUND [Tol, 1997], MERGE [Manne et al., 1995, Manne and Richels, 1992], PAGE [Hope, 2006] and WITCH [Bosetti et al., 2007]. All of these utilize optimal economic growth models to optimize the trade off between abatement today and climate damages in an uncertain future [Keller et al., 2004].

One major component of climate impacts that impedes consensus is the risk and uncertainty of climate catastrophes. Recent literature points out that neglecting climate risk can present seriously misleading optimal abatement policies [Pindyck, 2013, Kaufman, 2012, Ackerman et al., 2010, Weitzman, 2009, 2011, Martin and Pindyck, 2015]. Pindyck [2013], in particular, argues that the single most important driver of the social cost of carbon is the possibility of a catastrophic climate outcome. Policy that calls for aggressive precautionary abatement would be thus non-credible without considering the possibility of climate induced

catastrophes.

Among catastrophic climate outcomes, those that are assumed to take place when atmospheric temperature increase exceeds a climate temperature threshold, or *tipping point*, has entered numerous IAMs. Policy relevant tipping elements that are extensively discussed include the Atlantic thermohaline circulation (THC) collapse, Boreal forest dieback and Greenland ice sheet (GIS) and Arctic summer sea-ice melting, to name a few [Lenton, 2012, Lenton et al., 2008]. Incorporating catastrophes such as the THC and GIS to optimal growth models is nevertheless technically challenging. The optimal policy must optimize between lost opportunities for economic growth and climate damages that now encompass the abruptness and irreversibility of tipping elements. 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 et al., 2004]. Including the work of Keller et al., a number of studies have integrated a stochastic framework into IAMs such as DICE, to demonstrate the policy response to low-probability-high-impact catastrophic outcomes. These studies include Leach [2007], Manne and Richels [1992], Ackerman et al. [2010], Cai et al. [2012], Hope [2006], Lemoine and Traeger [2010, 2014], Keller et al. [2004], Webster [2002], Martin and Pindyck [2015], Dietz and Stern [2015], Weitzman [2009, 2011], Kelly and Kolstad [1999].

The development of climate IAMs reflect the improved scientific understanding of the climate system and increased attention to modeling uncertainty. As a result climate models have become larger and more complex, yet uncertainty of the climate system still remains predominantly unresolved [Urban et al., 2014, Pindyck, 2013]. Given potentially imminent risks of climate catastrophes coupled with slow rates of climate learning, the formulation of decision analysis tools to find room for agreement in near-term abatement has recently become an important task. Here I review the types of stochastic frameworks in the current

climate IAM literature that model decision making under uncertainty.

On one hand, efforts to realistically represent the complex structure of uncertainties and the integrated system restrict the embedding of stochasticity within the model and confine uncertainty analyses to parameterized decision rules, such as Monte Carlo simulations. Moles et al. [2004], Keller et al. [2004] have also noted that the non-convex nature of climate damages in large models can easily introduce computational problems such as non-smooth gradients and multiple local optima. In such cases, the use of commercially available optimization solvers to numerically solve fully stochastic model formulations may be limited. The most common, and sometimes the only available option these models have is to resort to Monte Carlo Averaging in assessing the uncertainty effect.

However, as Cai et al. [2012], Lemoine and Traeger [2010, 2014], Crost and Traeger [2013], Lemoine and Rudik [2016] point out, the Monte-Carlo averaging approach is misrepresentative of how we make decisions under uncertainty. The reason for this is that the approach is, in essence, an averaged sensitivity analyses, with each simulation equivalent to running a deterministic model [Crost and Traeger, 2013]. Despite their seeming *act-then-learn* nature of committing to a single policy under uncertainty, the many scenarios that the optimal policy averages over assume perfect information on the element of uncertainty. To plan from the beginning with knowledge on the time of occurrence and impact of a climate catastrophe, essentially takes away the most damaging aspect of uncertain climate outcomes—its *abruptness*. Monte Carlo Averaging not only outputs an optimal policy that is significantly underestimated, but can also wrongly conclude the sign of the uncertainty effect.

On the other hand, efforts to create a fully stochastic formulation of abrupt, irreversible climate outcome, along with an optimal policy in response to it, have steered the IAM community to integrate recursive dynamic programming (DP) frameworks to climate-economy

IAMs. Unlike deterministic models, stochastic DP frameworks allow for the model to treat tipping points as intrinsic random events that remain uncertain until catastrophic events are realized. Despite methodological advances in employing dynamic programming to climate policy however [Cai et al., 2012, Webster et al., 2012], conventional dynamic programming is not the most accessible of methods to obtain numerical solutions with; implementation is indeed computationally burdensome and dimensionally restricted, especially for assessing sequential decision policies characterized by a large set of actions [Powell, 2007, Sinha, 2005].

This means that for slightly larger climate-economy models such as DICE-WAIS [Diaz and Keller, 2016] that expands on the DICE model to account for a more detailed representation of sea-level rise and ice-sheet melting, there is a steep price to pay in terms of tractability for the level of solution accuracy recursive optimization methods are able to provide. And regarding the notion of *accuracy*—humility should be called for as it seems unlikely that we can obtain sufficiently precise estimates of uncertainty to settle on a $40 or an $80 carbon tax today. The current literature suggests that extant models are underprepared for a discussion on solution accuracy in stochastic IAM problems; perhaps to the extent that innovative dynamic programming implementation efforts are rendered somewhat unproductive.

For instance, inadequate representations of individual tipping elements in Global Climate Models (GCM) [Lenton, 2011, Pindyck, 2013]; the increased complexity of decision making under multiple (potential) tipping elements [Lenton et al., 2008, Lemoine and Traeger, 2016b, Martin and Pindyck, 2015]; wide disagreement in experts' judgement on probabilistic climate outcomes (*see figure 1.*, Kriegler et al. [2009], Zickfeld et al. [2007]); let alone the large ambiguity in the distribution of risk, all accomplish little for the added machinery of recursive methods, which otherwise would provide reliable *long-term* contingency plans for the stochastic problem. Given such uncertainties, a tentative solution should instead be intrin-

sically *operational*—an economical solution should encompass a set of robust, non-myopic *near-term* decisions that allow for sufficient expansion to the scope of both mitigatory actions and characterization of climate and economy [Brandimarte, 2014].

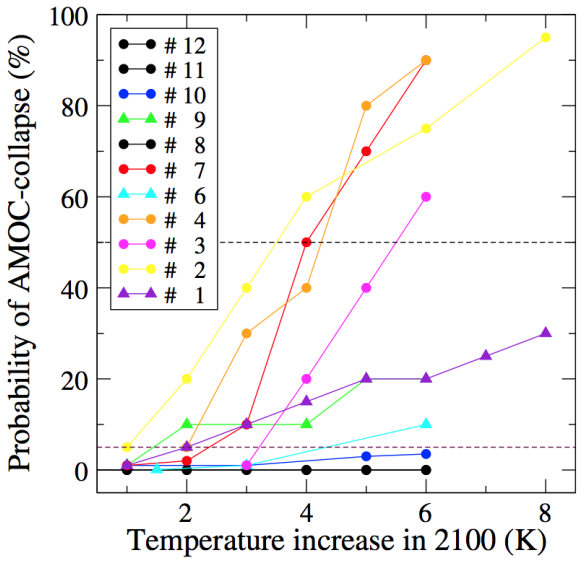

Figure 1. Differing Experts' Subjective Probability of THC Collapse as a Function of Temperature at Year 2100 Zickfeld et al. [2007]

In the following section, we provide a minimalist framework of precautionary mitigation using a *stochastic control* version of Nordhaus' DICE (2013), *DICESC*, to investigate the impact of uncertain catastrophic loss on near-term abatement. The model formulation encompasses the following: the stochastic program's recourse tree represents a dynamic cycle of decision making (level of emissions) and observational learning (state of climate); the tree also embodies abruptness of climate tipping (intrinsic uncertainty); following Lemoine and Traeger [2014], Kelly and Kolstad [1999], the tipping prior (or hazard rate) at each stage is endogenously determined by means of bayesian threshold learning; DICESC is conveniently written as a mathematical program and maintains a recourse structure that avoids the scaling limitations of classical multistage stochastic programming. We further stress the computational transparency of the model as DICESC is formulated as a one-shot optimization problem that

can replicate insights from recursive methods presented in Lemoine and Traeger [2014] and Cai et al. [2012], in a matter of minutes.

The paper outlines factors that contribute to a positive uncertainty effect on emissions abatement. Hedging for the risk of climate thresholds leads near-term abatement levels to increase, driven largely by precautionary incentives to decrease the future likelihood of climate tipping. We lastly assess the expected value of perfect information (EVPI)—the expected value of obtaining perfect information on the temperature threshold—to find that the EVPI can range anywhere from 0.4 to 5% of gross world output.

Sensitivity analyses on the distribution of risk however suggest that both the level of precautionary abatement and EVPI depend on higher moments of the Bayesian prior, as two priors with the same mean and variance are shown to display contrasting results. This exposes fundamental limitations of recommending robust mitigatory policies under uncertainty. The transparency of the model stresses that optimal prescriptions based on the most minimal machinery can have difficulty agreeing on magnitude, let alone the shape of the optimal control trajectory when knowledge of the tipping point probability distribution is insufficient.

2. DICESC: The Model

2.1 Recourse

DICESC adopts the DICE2013 model [Nordhaus and Sztorc, 2013, 2014] to demonstrate the relevance of stochastic control. DICE2013 (hereafter cited as DICE) is a recent version of the original climate IAM authored by Nordhaus in 1994 [Nordhaus, 1994] and calibrated to be consistent with the IPCC Fifth Assessment Report [Stocker et al., 2013]. This section presents modifications made to DICE. For readers who are unacquainted with the model, we state the main equations of interest in the Appendix section. While maintaining the simple nonlinear

program format of DICE, we introduce a recourse tree that integrates the various climate outcomes under requirement of *nonanticipativity*.

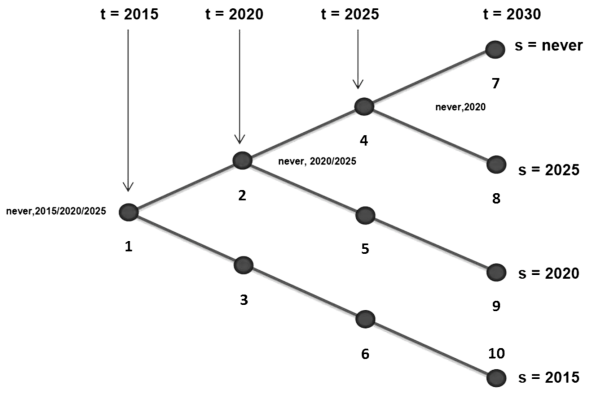

Figure 2.1. Stochastic Control Recourse Tree

To discretize the state space in stochastic programming, DICESC defines the set of periodic events with respect to occurrence of tipping. The time horizon of economic activity, t, 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 we add *states of the world (SoW)*, s, to represent scenarios associated with the year in which threshold damages take place. *Figure 2.1.* shows a simple recourse tree made up of four states of the world. For instance, $s = 2015$ represents the unlikely scenario that an irreversible climate catastrophe occurs at year 2015, which the planner only realizes in the subsequent period, 2020.

The *never* state naturally represents a world in which tipping does *not* take place and corresponds to the recourse path $[1 \rightarrow 2 \rightarrow 4 \rightarrow 7]$. Then for SoWs 2020 and 2025, the corresponding paths initially share the *never* path but deviate down the catastrophe branch at time 2020 and 2025, respectively. In each SoW, the binary parameter d_{cat}, in equation $(A5^{'})^{*}$ takes value 1 at $t = s$, to model an irreversible catastrophic outcome that takes place at the corresponding time. As a result of recourse, a state varying temperature threshold is applied

*see appendix

to damages $D(t)$ in expression $(A5)$ to output a state varying damage equation (1).

$$D(s,t) = \Psi_1 T^{AT}(s,t) + \Psi_1 [T^{AT}(s,t)]^2 + d_{cat} \cdot 2\left[.00644\left(\frac{T^{AT}(s,t)}{\bar{T}^{AT}(s)}\right)^3\right] \tag{1}$$

Note that by discretizing time, the temperature of the tipping threshold can be treated as a continuous variable. The corresponding tipping thresholds that are realized for each s in this setting, are equivalent to the temperature values that initiate tipping at $t = s$.

$$\bar{T}^{AT}(s) = T^{AT}_{t=s}(never) \tag{2}$$

2.2 Terminal Condition

To approximate choices over the infinite horizon, we use a sequential state variable targeting approach described in Lau et al. [2002]. It is assumed that economic activity is deterministic in the post-economic time horizon; that is, the climate time horizon is driven by the *post-terminal growth factor*, γ', which is defined by the exogenous growth rates of population and total factor productivity;

$$\gamma'_t = \left(\frac{L_t}{L_T}\right)\left(\frac{A_t}{A_T}\right)^{\frac{1}{(1-\gamma)}} \tag{3}$$

where γ is the elasticity of capital in the production function, and T denotes the terminal time period of the economic time horizon. The target state variable of interest is the *post-terminal capital stock*, $\bar{K}_{T+1}$, which in turn determines the trajectory of the control variable, investment (I_t). Hence, by iteratively updating the value of $\bar{K}_{T+1}$ using the post-terminal growth factor;

i.e.

$$\bar{K}_{T+1}(\texttt{iter}) = \gamma_t' \cdot K_T(\texttt{iter}) \tag{ITR}$$

and subsequently solving for a smooth investment trajectory, I_t, $t = 2015, \ldots, T$ by

$$\bar{K}_{T+1} = (1-\delta)K_T + I_T \tag{TC}$$

$$K_{t+1} = (1-\delta)K_t + I_t$$

we obtain a set of control trajectories that mimic a steady state economy in the post-terminal period. DICESC updates the terminal capital stock iteratively using a sequential solve of the model until convergence in the investment trajectory is achieved throughout the time horizon. Convergence is achieved within 20 solve iterations. Notice that this approach does not require the model to achieve the steady state growth rate by period T, nor does it impose a target capital stock in the post-terminal period.

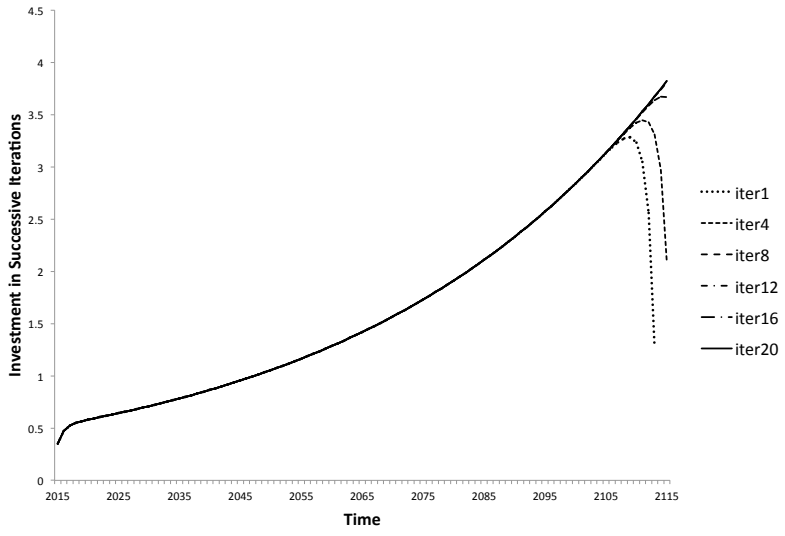

Figure 2.2. Terminal Condition: Iterative State Variable Targeting

3. Stochastic Control

3.1 Stochastic Programming with Recourse: Act Then Learn

We have so far introduced the concept of recourse to DICE. In this section we demonstrate how this recourse tree is used to output *feasible* policies that address the risk of threshold effects. For policies to be implementable, decisions must be a function of the history of the event process available at the time decisions are made. This requirement is called the *nonanticipativity constraint*, which states that at every time period, decisions are based only on available information. To formalize the concept of nonanticipativity, we follow the notation used in Shapiro [2009].

Let Ω_t denote the set of all nodes at stage $t = 2015, 2020, \ldots, T$. $\xi_t \in \Omega_t$ is then the state of the system at time t, which is defined by state of the world s. The *children nodes* of node ξ_t correspond to the nodes that can take place given ξ_t at $t+1, t+2, \ldots, T$. For $1 \leq s \leq t \leq T$, denote $\xi_{[s,t]} := \xi_s, \ldots, \xi_t$, which represents the history of the process from node ξ_s to ξ_t. Feasible decision policy $x_t(\xi_{[2015,t]}) \in \mathbb{R}^{n_t}$ is adapted to the time structure of the process described by the following *act-then learn* process:

$$\text{act}(x_{2015}) \longrightarrow \text{learn}(\xi_{2015}) \longrightarrow \text{act}(x_{2020}) \longrightarrow \text{learn}(\xi_{2020}) \longrightarrow \ldots \tag{4}$$

Decision vectors that satisfy the feasibility (nonanticipativity) constraint can accordingly be written as:

$$x_{2015} \in \chi_{2015}, \;\; x_t \in \chi_t(x_{t-1}, \xi_t), \quad \text{for } t = 2020, \ldots, T \tag{5}$$

where $\chi_{2015} \subset \mathbb{R}^{n_1}$ is a deterministic set and $\chi_t : \mathbb{R}^{n_{t-1}} \times \Omega_t \rightrightarrows \mathbb{R}^{n_t}$ is a point-to-set mapping.

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 will experience threshold damages. Let us consider a constant probability density of a hazard $hr_t = Pr(\xi_{t+1}^{t+1} | \xi_t^s = \xi^{never})^*$. For each time period, the planner is able to hedge against the risk of tipping in the immediate next period given the rate of hazard. In the next period the planner realizes one of two states $\xi_{t+1}^s = \left[\xi^{never}, \xi^{t+1}\right]$; the planner either faces a state of threshold damages, or finds itself continuing down the never path, yet to encounter the problem of hedging against the next period's uncertain outcome.

Given hr_t, the conditional probability of climate tipping during period t, assuming that there has been no tipping to that point in time, is used to assign probability Π_s, for each SoW path. That is,

$$\begin{aligned} \Pi_s &= sr_t \cdot hr_t, \quad t = s \\ sr_t &= \prod_{time=0}^{t} (1 - hr_t) \end{aligned} \tag{6}$$

where sr_t is the survival rate, the probability of *no climate tipping* until time $t = s$. Each s is hence treated as a *deterministic* run that makes up the optimal solution to the stochastic problem. At this point however, I note that solving for each deterministic SoW path in DICESC is in essence, *not* fully deterministic.

When perfect information about a uncertainty scenario is assumed, each deterministic run becomes a *learn-then act* optimization problem. This problem formulation is easily implemented, but limits adequate model representation of sequential decision making under uncertainty, especially when recourse options are available. The subtle inadequacy of rep-

*ξ_t^s denotes the realized state, s, at time t.

resenting uncertainty in *learn-then act* form relates to the assumption that all uncertainty is unraveled before the model is run. The concept of nonanticipativity has no place in this setting, as perfect information regarding the entire sequence of events is known in the initial time period. Under such heroic assumptions, the economy will plan accordingly to output a smooth consumption path, essentially eliminating the impact of *abruptness* when uncertain catastrophic outcomes are suddenly realized.

The appropriate deterministic setting we should be imposing then, concerns one that embodies *act-then learn*, so that despite the pre-determined time of tipping at time t, adaptive measures can only take effect after tipping is realized; i.e., information on state ξ_t arrives. Such sequential resolution of uncertainty by means of learning is best depicted in the advancement of IPCC Assessment Reports. The IPCC Fifth Assessment Report (AR5) states:

> "*The estimates of sea-level rise presented in AR4 explicitly excluded future rapid dynamical changes in ice flow, and stated that "understanding of these processes is limited and there is no consensus on their magnitude". Considerable efforts have been made since AR4 to fill this knowledge gap... The number of in situ and satellite observations of cryospheric parameters has increased considerably since AR4 and the use of the new data in trend analyses, and also in process studies, has enabled increased confidence in the quantification of most of the changes.*"

As technological advancement is evidently used to understand the process of climate change with *higher confidence*, learning achieved *post* AR4 is used to re-diagnose the current state of climate in AR5. With respect to the risk of large scale climate outcomes, Smith et al. [2009] confirms that

> "*there is now better understanding that the risk of additional contributions to sea level rise from melting of both the Greenland and possibly Antarctic ice sheets may be larger than projected by ice sheet models assessed in the AR4...*"

This course is in a sense, compatible to a periodic evaluation of hazard rates, which makes act-then learn increasingly relevant to modeling any uncertain climate outcome. *Figure 3.1.a & 3.1.b.* provides a comparison between the optimal abatement paths under learn-then act

and act-then learn. In learn-then act, the planner plans in full knowledge of the threshold at which tipping takes place, as well as the potential climate damages ascribed to the *time* of tipping. In contrast, act-then learn assumes that the climate threshold remains unknown until tipping occurs, and is realized once threshold damages take effect. At year 2030, the planner learns whether or not tipping has occurred in the previous period; *Figure 3.1.b.* hence represents hedging strategies that diverge based on different realized climate outcomes.

Figure 3.2 shows the recourse structure of optimal actions based on climate scenarios in an act-then learn stochastic programming format. Note that higher levels of abatement along the act-then learn *never* path illustrate precautionary incentives for abatement under abrupt climate catastrophes. To sum up, by assigning a probability to each *act-then learn* scenario path, the objective of the model is to engage in a comprehensive hedging strategy against tipping uncertainties throughout the stochastic horizon.

3.2 Stochastic Control: Endogenous Learning & Controlled Hazard Rates

In a stochastic programming (SP) framework, the hazard rates, and thus the scenario probabilities, are exogenously determined. Here, the economy executes a relatively passive level of learning. While the economy could apply periodic updates to the tipping point distribution as it progresses down the *never* path, stochastic programming does not allow for the hazard rate to be updated accordingly, so learning remains ineffective. Lemoine and Traeger [2014] establishes the role of Bayesian learning about the location of the threshold. *Learning* here refers to *learning from experience* or *learning by observation*, discussed in Kelly and Kolstad [1999]. Unlike *passive* or *autonomous* learning for which learning corresponds to the exogenous arrival of information [Webster, 2002], learning from observation entails experimenting with emissions to obtain information about the tipping point. *Endogenous learning* is then characterized as the continual process of establishing the economy's state of knowledge by observing the climate,

and optimizing actions with the information at hand. Following Lemoine and Traeger [2014], I adopt a simple prior; for tipping points $\bar{T}^{AT}$, and assume that the planner is fully aware of the distribution. Note that the maximum temperature threshold in the domain, $\bar{T}_{max}$, is the temperature at which the planner believes with certainty that climate tipping occurs.

$$\bar{T}^{AT} \in [\bar{T}_0, \bar{T}_{max}] \tag{7}$$

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:

$$\bar{T}^{AT} \in [T_t^{AT}(never), \bar{T}_{max}] \tag{8}$$

A more active type of learning, learning through observation, is thus incorporated into the hazard rate:

$$hr_t(x) = \frac{\Phi_{\bar{T}^{AT}}[T_{t+1}^{AT}(x)] - \Phi_{\bar{T}^{AT}}[T_t^{AT}(x)]}{\Phi_{\bar{T}^{AT}}[\bar{T}_{max}] - \Phi_{\bar{T}^{AT}}[T_t^{AT}(x)]} \tag{9}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the known Bayesian prior and x, the emissions abatement trajectory over time. Note that the endogeneity of hazard rate stems from the global temperature's dependence on abatement trajector, x.

When assuming a uniform prior as in Lemoine and Traeger [2014], the hazard rate expression in (9) is equivalent to:

$$hr_t(x) = \frac{T_{t+1}^{AT}(x) - T_t^{AT}(x)}{\bar{T}_{max} - T_t^{AT}(x)} \tag{10}$$

Based on the hazard function, a rise in temperature incurred between times $[t, t+1]$ increases

the probability of experiencing tipping at the end of the time period. The abatement level at time t hence influences the likelihood of risky outcomes tomorrow and alters the probabilities assigned to each branch (SoW) of the recourse tree. While stochastic programming could only justify mitigatory abatement to minimize expected threshold damages, stochastic control provides reason to abate today as to change the likelihood of the occurrence of the risky outcome. Note also that there exists a trade off between resolution of uncertainty and the benefits of mitigatory efforts because GHG control decreases the rate of learning [Kelly and Kolstad, 1999].

3.3 Stochastic Control vs. Stochastic Programming

A high-level problem formulation for stochastic control and stochastic programming can be expressed as a function of abatement.

Stochastic Programming

$$\max_{x} \sum_{s} \Pi_s \cdot F_s(x) \tag{11}$$

Equation (11) characterizes the stochastic programming formulation in which Π_s denotes the probability assigned to state of the world, and x, the abatement trajectory vector for the economic time horizon. $F_s(x)$ is the Lagrangian (the model) indexed by state of the world that represents the cost and benefits of emissions abatement. Note that abatement vector x only enters the problem through state indexed Lagrangians. In the stochastic programming formulation, abatement vector x satisfies the following first order condition:

$$\underbrace{\sum_{s} \Pi_s \frac{dF(x)}{dx}}_{\text{damage minimizing incentive}} = 0 \tag{12}$$

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

Stochastic Control

$$\max_{x} \sum_{s} \Pi_s(x) \cdot F_s(x) \;\; s.t. \tag{13}$$

$$\Pi_s(x) = sr_t(x) \cdot hr_t(x), \quad t = s$$

$$hr_t(x) = \frac{\Phi_{\bar{T}_{AT}}[T^{AT}_{t+1}(x)] - \Phi_{\bar{T}_{AT}}[T^{AT}_{t}(x)]}{\Phi_{\bar{T}_{AT}}[\bar{T}_{max}] - \Phi_{\bar{T}_{AT}}[T^{AT}_{t}(x)]}$$

$$sr_t(x) = \prod_{time=0}^{t} (1 - hr_{time}(x))$$

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

$$\underbrace{\sum_{s} \Pi_s(x) \frac{\partial F_s(x)}{\partial x}}_{\text{damage minimizing incentive}} + \underbrace{F_{never}(x) \frac{\partial \Pi_{never}(x)}{\partial x} + \sum_{s^{-never}} F_s \frac{\partial \Pi_s(x)}{\partial x}}_{\text{precautionary incentive}} = 0 \tag{14}$$

Equation (14) shows that optimal abatement in a stochastic control setup is motivated by two incentives; an incentive to minimize threshold damages given probabilistic risky outcomes, and furthermore a precautionary incentive to delay or even avoid such catastrophes by means of hazard rate control. Policy optimizes the likelihood of states such that damage minimization results in the highest possible welfare.

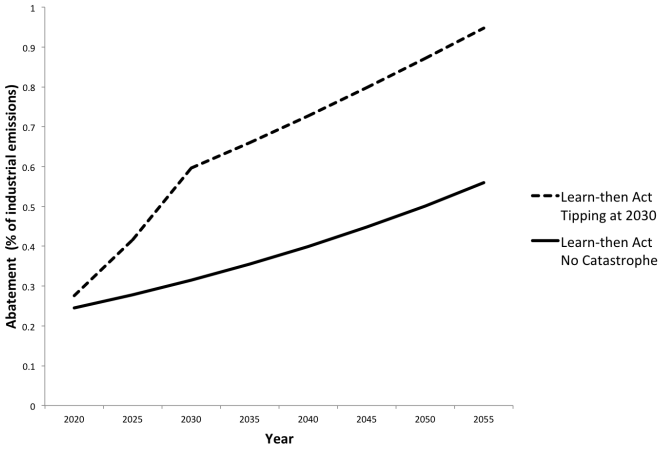

Figure 3.1.a. Learn-then Act

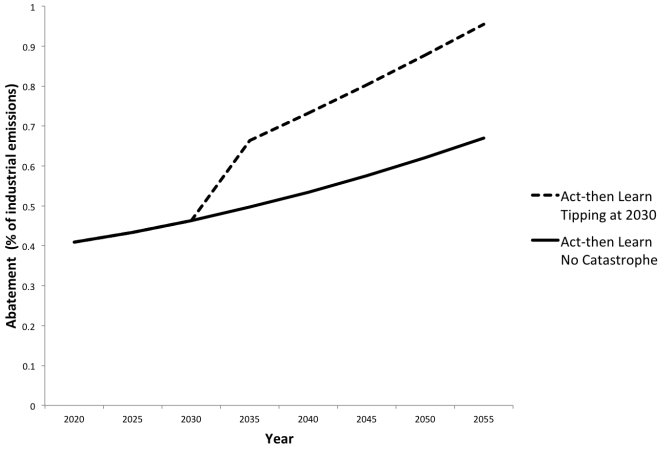

Figure 3.1.b. Act-then Learn

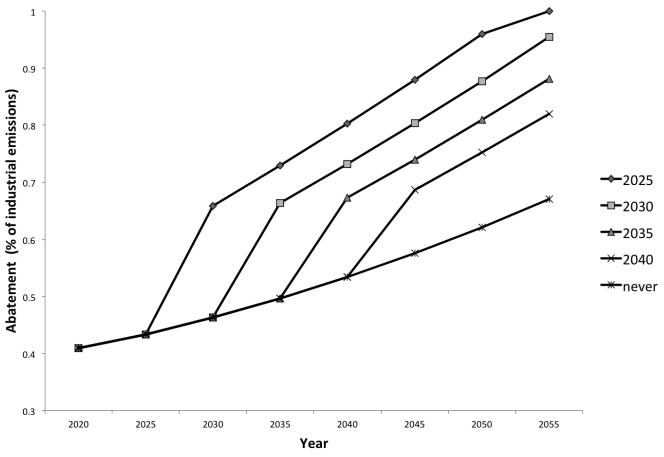

Figure 3.2. Optimal Near-term Abatement in Act-then Learn

4. Optimal Carbon Abatement Trajectory

This section illustrates near-term abatement trajectories obtained by DICESC when hedging against the risk of irreversible climate tipping with uncertain thresholds. Following Stocker et al. [2013], Lenton et al. [2008], we assume at least one tipping element is triggered with certainty when global temperatures exceed the expected equilibrium climate sensitivity of 3 °C, and model for a single tipping element with an unknown temperature threshold. This assumption is also in accordance with Kriegler et al. [2009], which deduces conservative lower bounds to the probability of triggering at least one large-scale tipping point, including the Greenland and West Antarctic ice sheets, Atlantic meridional overturning circulation (AMOC) and the Amazon rainforest. Among policy relevant tipping elements, *Arctic ice-sheet melting (ASI)* is an adequate instance to model using DICESC due to the its short transition time and associated uncertainty domain of the tipping threshold.

The Arctic sea-ice melting (ASI) is characterized to be set off by an increase in global atmospheric temperature between $[0.5, 2]$ °C. ASI's transition timescale to a new stable equilibrium is projected to be only 10 years; among the policy-relevant tipping elements discussed in Lenton et al. [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 in contrast 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, since the absence of ice exposes a darker ocean that absorbs more radiation. Lenton et al. even suggests 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, along with DICESC's assumption that damages are realized in the subsequent post-tipping period, the relatively short transition timescale of the tipping element makes Arctic sea-ice melting an

adequate tipping element to model.

The stochastic horizon is fixed to be 80 years, making the last tipping state of the world to experience threshold damages starting at year 2095. Lastly, we define $\bar{T}^{AT}$ as the increase in global atmospheric temperature since year 1900 and set the prior distribution domain for the temperature increase threshold to be $\bar{T}^{AT} \in [\bar{T}_0, \bar{T}_{max}] = [0.83, 3.0]\,^{\circ}\text{C}$. 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 run time for DICESC with $\bar{T}_{max} = 3.0$ with a stochastic horizon of 80 years is under 3 minutes using NLP solver *CONOPT* in GAMS on a laptop computer.

4.1 Optimal Abatement

This section compares three optimal abatement policy scenarios.

- **DICE2013**: Optimal policy in the original DICE model
- **DICESP**: DICE stochastic programming with fixed 5% hazard rate
- **DICESC**: DICE stochastic control with $\bar{T}_{max} = 3.0\,^{\circ}\text{C}$

DICE2013 consists of the original version of DICE that does not assume climate tipping. For stochastic control, we assume a normal distribution for the tipping points symmetrically distributed on $[0.83, 3.0]$ (*figure 6.*. The hazard rate that exogenous *SoW* probabilities in stochastic programming, is fixed at 5%. To provide a benchmark, Lontzek et al. [2012]'s calculations based on expert probability estimates of THC collapse (estimated to be triggered by $+[3, 5]\,^{\circ}\text{C}$) in Zickfeld et al. [2007] conclude an $\approx 7.5\%$ hazard rate given global warming of $3.0\,^{\circ}\text{C}$, estimated by an average pessimistic expert. The corresponding hazard rate estimated by an average optimistic expert is $\approx 2\%$. Comparing policies resulting from stochastic programming and stochastic control allows for a first order decomposition of abatement that

exposes precautionary abatement—additional abatement that arises due to the endogeneity of the hazard function. Lemoine and Traeger [2014] shows the decomposition of optimal abatement using first order derivatives similar in nature to equations $(19) - (21)$, and denotes the contribution of endogenous hazard rates, the *marginal hazard effect*. We replicate this decomposition numerically in *Figure 4.*

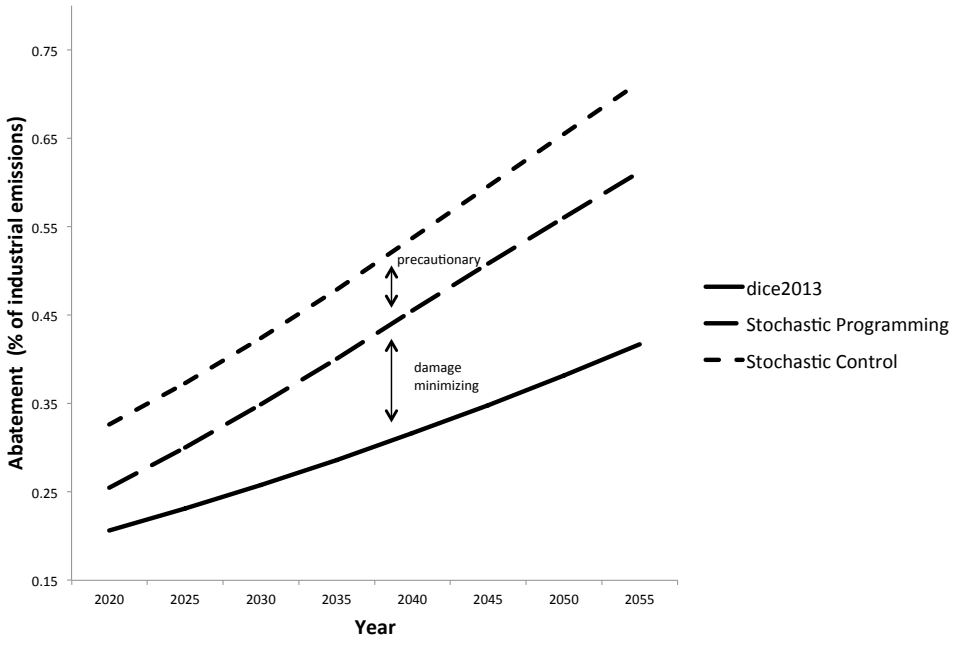

Figure 4. Optimal Near-Term Abatement

The difference in abatement levels between DICESP and DICE2013 alone shows the increase in near-term abatement when positive probability of a high risk outcome is introduced to the model. While this discrepancy arises due to incentives to minimize expected threshold damages, it does not exhibit precautionary abatement since emissions control cannot help delay or avoid threshold damages. In a stochastic control setting, emissions abatement can further be used to alter hazard rates, and as a result policy commits significantly more to near-term abatement. *Figure 4.* displays the portion of optimal abatement used to *optimize* hazard

rates; i.e. to determine the optimal balance between consumption and *expected catastrophic loss*.

4.2 Considering Alternative Priors: Mean Preserving Spreads (MPS)

In this section, we perform sensitivity analyses on the distribution of tipping points. The initial prior employed in *section 4.1* is a normal distribution calibrated so that the upper and lower tail decline to $\epsilon > 0$ at each end of the domain, 0.83 °C and 3.0 °C, respectively. Note that, in this distribution a tipping threshold of 1.915 °C is assigned the highest likelihood, with equally insignificant weights assigned to the worst and best case scenarios. Making this our baseline model, we follow Baker [2009] to apply the notion of *increasing risk* defined by Rothschild and Stiglitz [1970] to the baseline prior.

Definition. *Z is risker (or a Mean Preserving Spread (MPS)) than Z′ iff* $\mathbb{E}_Z U(Z) \leq \mathbb{E}_{Z'} U(Z')$

An obvious choice for a mean preserving spread is a uniform prior (denoted *MPS: Uniform*), which results in a considerably lower objective value for all maximum tipping thresholds considered (*figure 6.*). This is because the decision maker recognizes that *all* tipping points in the domain, including ones in the left end of the tail, are equally likely to trigger a catastrophe when assuming a uniform tipping point distribution. we keep the first two moments of the uniform distribution constant and create a second MPS; a truncated normal distribution (denoted *MPS: Truncated Normal*) on temperature domain $[0.83, \infty]$. *Figure 5* plots the three priors used in the current analysis.

While all trajectories display a robust incentive for precautionary abatement, optimal near-term abatement levels for each MPS prior (displayed in *figure 7.*) show a notable increase relative to that of the baseline, especially in the immediate future. This is caused by the transfer of weight to the lower end of the tipping distribution characterized by adverse states

of the world in which tipping is triggered by small increases in atmospheric temperature. Lastly, recall that the two MPS priors (truncated normal and uniform priors) are calibrated to share the first and second moments of their respective distributions. Results indicate that the uniform prior brings about an abatement trajectory that is approximately 8% greater in the initial period than that of the truncated normal MPS. The difference gradually decreases to 2% in the year 2055. Although the difference between the policies is moderate, it is clear that near-term optimal abatement not only depends on the mean and variance of risk, but also on the higher moments of the Bayesian prior.

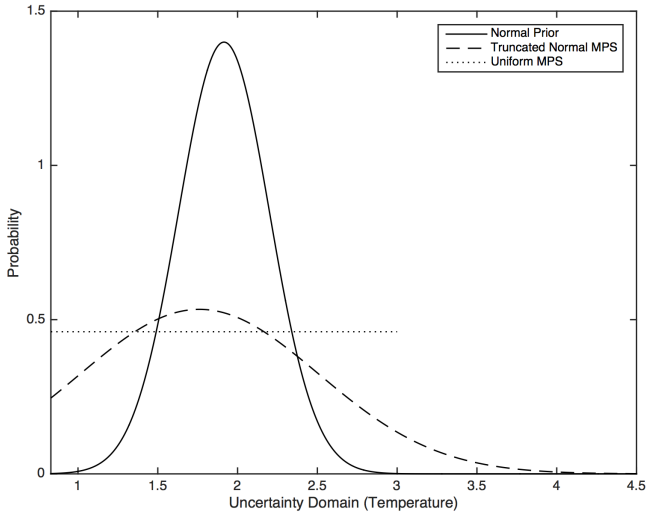

Figure 5. Tipping Prior Distributions

5. Expected Value of Perfect Information

The discussion thus far brings to light key factors that determine near-term optimal abatement under tipping point uncertainty; the possibility of triggering an adverse climate outcome, the responsiveness of hazard rates to the economy's mitigatory efforts and lastly, the assumed distribution of risk. In this section we show that the higher moments of the assumed risk prior is a critical component in determining the first order magnitude of the *Expected Value of*

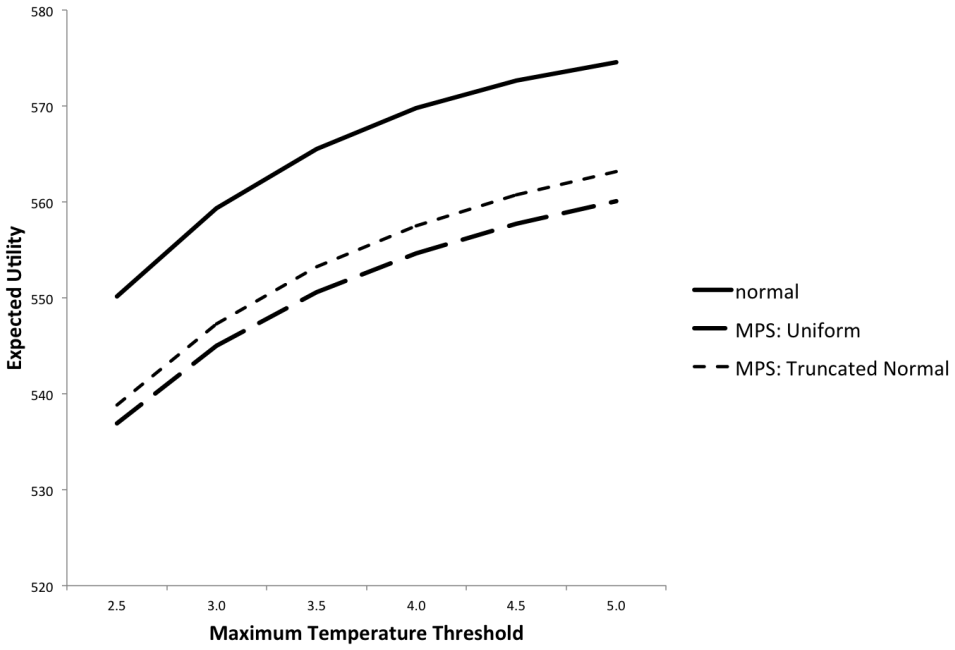

Figure 6. Expected Welfare Across Tipping Priors

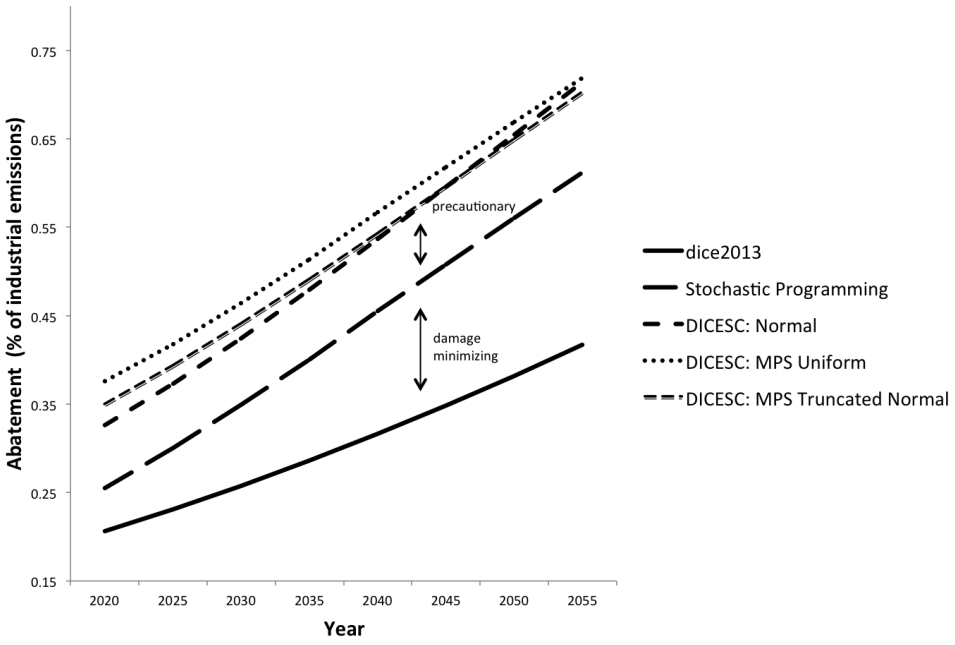

Figure 7. Optimal Near-Term Abatement Across MPS

Perfect Information (EVPI)—the value of resolving all uncertainty to obtain perfect information of the temperature threshold. We adopt the approach presented in Manne and Richels [1992] to compute the EVPI as the difference between the expected welfare resulting from optimization under uncertainty, and the expected welfare under perfect information. Denote $\mathcal{P}$ to be the assumed prior, and $f^{\mathcal{P}}$ to be the probability density function that corresponds to the respective priors. Also denote $U^*_{\bar{T}^{AT}}$ as the highest achievable utility given perfect information on the tipping point being equal to $\bar{T}^{AT}$, and $U^{\mathcal{P}}_{\bar{T}^{max}}$ as the expected utility output by DICESC given prior $\mathcal{P}$ and the domain of the distribution, $\bar{T}^{max}$ (note that $\bar{T}^{max} = +\infty$ for the truncated normal distribution).

$$EVPI^{\mathcal{P}}_{\bar{T}_{max}} = \underbrace{\int_{\bar{T}_0}^{\bar{T}_{max}} U^*_{\bar{T}^{AT}} f^{\mathcal{P}}(\bar{T}^{AT}) d\bar{T}^{AT}}_{deterministic} - \underbrace{U^{\mathcal{P}}_{\bar{T}_{max}}}_{stochastic} \tag{15}$$

EVPI results are displayed in *Table 1.* While the EVPI of the normal and uniform distributions are both under 1% of gross world output, the value based on the truncated normal prior is over 5%.

	Welfare		EVPI	
Prior	Deterministic	Stochastic	Value	Percentage
Normal	562.234	559.590	2.644	0.472%
MPS: Uniform	549.338	547.368	1.970	0.360%
MPS: Truncated Normal	583.692	553.737	29.955	5.410%

Units are in Gross world output (trillions 2005 USD)

Table 1. EVPI Across Priors

This discrepancy in EVPI stems from the weight assigned to the right-end of the tail (above 3 °C) in the truncated normal MPS. Since there is now a non-trivial chance that the tipping

point is beyond 3 °C when assuming the truncated normal MPS, the *expected utility under perfect information* (denoted *deterministic*) is considerably greater than that of the remaining priors. In other words, added weight on scenarios in which the tipping point is not exceeded contributes to a high learn-then act outcome. This however does not mean that the truncated normal MPS also outputs high expected welfare under uncertainty. On the contrary, *Figure 6.* shows a steep drop in expected welfare when applying the truncated normal MPS to the baseline prior. It is clear that the substantial weight increase in the left-end of the tail (adverse scenarios in which climate tipping takes place early) forces the act-then learn decision maker to optimize the tradeoff between abatement today and *riskier* outcomes in the future, resulting in heightened precautionary abatement and lower expected welfare.

Recall once again that the two MPS priors have the same mean and variance. EVPI analyses however show that two distributions with the same first and second moments may argue for very different EVPIs. Information on the higher moments of the Bayesian prior are thus critical in determining the value of information.

6. Conclusions

With climate change, learning is continuous. The series of IPCC reports serve as an indicator for steady advancements in observational learning—improved satellite sensing technologies and on-site measurement techniques—have progressively filled the knowledge gap of climate uncertainty. A Bayesian approach to learning in an *act-then learn* framework is hence powerful, as it captures the sequential process of decision making under uncertainty based on *new observational evidence*. In an act-then learn stochastic control framework, we show that the possibility of climate tipping in the future considerably increases optimal abatement to delay or even avoid the occurrence of threshold damages.

Distributional sensitivity analyses nevertheless expose the stochastic framework's strong

dependency on higher moments of the Bayesian prior. A decomposed analysis of abatement reveals that precautionary abatement incentives are sensitive to higher moment distributive properties. The dependency on higher moments is more severe when assessing the expected value of information. This is discomforting as it seems unlikely that we can determine the distribution of risk with sufficient precision. Decision-dependent uncertainty, or stochastic control, under an act-then learn framework is no doubt fundamental in modeling decision making under climate uncertainty. Nevertheless it is evident that the approach suffers from the pitfalls in Bayesian learning, even in the most minimal frameworks.

Appendix to Chapter 2: DICE 2013, Nordhaus

$$W = \sum_{t=1}^{Tmax} \frac{1}{(1+\rho)^t} U[c(t), L(t)] \tag{A1}$$

$$U[c(t), L(t)] = L(t)\left[\frac{c(t)^{1-\alpha}}{(1-\alpha)}\right] \tag{A2}$$

$$E_{Ind}(t) = \sigma(t)[1-\mu(t)]Q(t) \tag{A3}$$

W in equation (A1) 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 corresponds to 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 (A3) describes the amount of GHG emitted through economic production, $Q(t)$. Here $\sigma(t)$ denotes the GHG emissions to output ratio and $\mu(t)$, the emissions-reduction rate (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 control variable central to this research.

$$Q(t) = \frac{[1-\Lambda(t)]A(t)K(t)^{\gamma}L(t)^{1-\gamma}}{[1+\Omega(t)]} \tag{A4}$$

$$D(t) = \Psi_1 T^{AT}(t) + \Psi_1 [T^{AT}(t)]^2 \tag{A5}$$

$$\Lambda(t) = \theta_1(t)\mu(t)^{\theta_2} \tag{A6}$$

Equations (A4)-(A6) express the dynamics of climate damages and costs as a fraction of gross output. Q denotes economic output net of climate damages ($D(t)$) and abatement costs ($\Lambda(t)$), where $A(t)$ is the Hicks-neutral total factor productivity and $K(t)$, the capital stock. In more detail, equation (A6) expresses the ratio of abatement costs to output, $\Lambda(t)$, as a convex function of abatement and (A5) presents damage impacts from climate change as an increasing function of mean global atmospheric temperature.

Estimation of the damage function has been controversial, even more so than the projections of climate change [Schneider, 2009]. Nordhaus and Sztorc [2013, 2014] point out that equation (A5) is a highly simplified damage function that takes into account both monetized damages (that rely on Tol [2009]) and non-monetized impacts such as loss of biodiversity, political reactions to climate change and impacts that are difficult to model (extreme events and long-term warming). Although the damage function is consistent with other major studies in the literature, Nordhaus remarks that the damage function ultimately does not consider tipping points.

To consider tipping elements, we use Nordhaus' DICE-2013 GAMS code presented in Nordhaus and Sztorc [2014] which includes a threshold damage term in the damage function (see equation $A5^{'}$).

$$D(t) = \Psi_1 T^{AT}(t) + \Psi_1 [T^{AT}(t)]^2 + d_{cat} \cdot 2\left[.00644\left(\frac{T^{AT}(t)}{\bar{T}^{AT}}\right)^3\right] \tag{A5'}$$

The last term in equation $(A5^{'})$ hence represents irreversible damages from climate tipping, which takes effect (turns on and stays on) when atmospheric temperature exceeds the threshold temperature, $\bar{T}^{AT}$. Note that this term can be switched on or off by means of the binary

parameter, d_{cat}.

$$M_{AT}(t) = E_{Ind}(t) + E_{Land}(t) + \phi_{11} M_{AT}(t-1) + \phi_{21} M_{UP}(t-1) \tag{A7}$$

$$F(t) = \eta log_2 \left[\frac{M_{AT}(t)}{M_{AT}(1750)} \right] + F_{EX}(t) \tag{A8}$$

$$T^{AT}(t) = T^{AT}(t-1) + \xi_1 F(t) - \xi_2 T^{AT}(t-1) - \xi_3 [T^{AT}(t-1) - T_{LO}(t-1)] \tag{A9}$$

The last set of equations describe the process of how industrial emissions increase the global mean atmospheric temperature. Equation (A7) shows that industrial emissions adds to the carbon concentration in the atmosphere $M_{AT}(t)$, which in turn increases the radiative forcing of greenhouse gases $F(t)$ (A8). Lastly and most importantly, equation (A9) 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 lacking empirical support. To settle on a concrete policy that requires minimal margin for error such as the social cost of carbon, a more detailed IAM should be used [Nordhaus, 2011]. However, the transparency of the model often comes in handy when introducing methodological steps to climate IAMs; those that suggest effective directions in which optimal policy should move toward, especially in discussing policy under uncertainty.

Chapter 3.

Optimal Climate Policy: The Importance of "Act-then Learn" Under the Risk of Climate Tipping

1. Introduction

Determining the effect of uncertainty in complex integrated assessment models (IAMs) has become increasingly relevant in the context of low-probability, high-impact climate catastrophes. Naturally, the discussion regarding the proper (and improper) frameworks to treat such an effect has emerged, resulting in the output of pedagogic papers and model documentations such as Lemoine and Rudik [2016], Chang and Rutherford [2017], Crost and Traeger [2013], Cai et al. [2012], Webster et al. [2012].

We can categorize stochastic frameworks employed by IAMs using the concept of *act-then learn* and *learn-then act* introduced in Manne and Richels [1992]. These regulatory strategies are differentiated by *when* learning takes place; i.e. the point in the decision making process at which uncertainty is resolved. In learn-then act, the decision maker is fully informed of the uncertain element of the model before policy is optimized. In constrast, act-then learn corresponds to the case in which the decision maker first commits to an action under uncertainty and learns of the uncertain outcome or state of the world at a future time period. Act-then learn is generally viewed as the proper strategy to treat uncertainty. Climate models based on stochastic dynamic programming [Kolstad, 1996, Cai et al., 2012, Lemoine and Traeger, 2014],

approximate dynamic programming [Webster et al., 2012] and stochastic control [Chang and Rutherford, 2017] all incorporate act-then learn.

The integrated assessment community however, has more often resorted to using Monte Carlo methods to produce an uncertainty effect [Hope, 2006, Dietz, 2011, Nordhaus, 2014, Ackerman et al., 2010, Keller et al., 2004, Diaz and Keller, 2016]. The most commonly used Monte Carlo method is Monte Carlo Averaging (MCA). MCA employs learn-then act—the many different selves of the decision maker are ex-ante given perfect information on the uncertain elements of the model, which makes each simulation equivalent to running a deterministic model [Crost and Traeger, 2013]. The optimal policy, x^*, resulting from an MCA framework can be written as follows:

$$x^* = \frac{1}{n}\sum_{s=1}^{n} x_s^* \qquad \text{(MCA)}$$

$$x_s^* = arg\max_{x \in X} \; F_s(x;\theta^s)$$

where F is the model and s, the state index for which vector of uncertain parameters, θ^s is drawn using Monte Carlo sampling. x_s^* thus represents the state-specific optimal policy given perfect information on uncertain parameters of the model.

Averaging policies output by MCA simulations is equivalent to averaging across states of the world that heroically assume all uncertainty is resolved before the initial period. MCA, or any type of learn-then act framework for this matter not only misrepresents how decisions are made under uncertainty, but also leads the resulting policy to be much less demanding than it should be under proper treatment of uncertainty.

However, note that not all Monte Carlo methods employ learn-then act. Keller et al. [2004] uses *Monte Carlo Fitting (MCF)*, an open-loop system strategy equivalent to applying sample

average approximation over Monte Carlo sampled states [Kleywegt et al., 2002, Homem-de Mello and Bayraksan, 2014]; that is, MCF identifies a single policy to optimize over the range of simulated outcomes. The optimal policy, x^* under MCF can be expressed as follows:

$$x^* = arg \max_{x \in X} \ \frac{1}{n} \sum_{s=1}^{n} F_s(x; \theta^s) \tag{MCF}$$

Technically, MCF simply employs *"act"*, with no scope for learning. The decision maker commits to an optimal policy trajectory that takes into account all scenarios under uncertainty. Such a policy nevertheless fails to take into account possible future realizations of the true state of the world. In the past, Arrow and Fisher [1974], Epstein [1980], and more recently, Chichilnisky and Heal [1998], Webster [2002], have focused on how the ability to learn in the future influences emissions policy today under the risk of uncertain and irreversible climate impacts. This small stream of research suggests that if the influence of learning to current policy is significant, then policies derived using MCF can be potentially misleading. In this paper, we demonstrate *how* misleading the resulting policy can be in the absence of *learning*.

We consider two types of learning, active and passive, as previously defined by Kolstad [1996]. Active learning refers to the generation of information through the control variable, which in our case, is carbon emissions. The process of Bayesian updating the prior distribution of risk is a form of such learning. Passive learning on the other hand, is more *observational* in nature and refers to the exogenous arrival of information every period. For instance, the decision maker's ability to keep track of the history of the event process, such as whether or not climate tipping has occurred, takes the form of passive learning. We show that the absence of sequential learning in either type can significantly undermine near-term precautionary abatement incentives, and argue that under the risk of *near-term* climate tipping, act-then learn is the only strategy that can replicate a reasonable uncertainty effect.

Note that the nature of our research bears resemblance to the work of Crost and Traeger [2013] on the shortcomings of MCA relative to frameworks that embody *decision making under uncertainty*. This paper takes a step further to show that decision making under uncertainty on its own may not suffice in prescribing a near-term optimal carbon abatement policy, especially when passive learning is a large part of the sequential learning process. The paper is organized as follows: in section two, we lay out the model (DICESC) used to represent act-then learn. Sections three and four present the lack-of-learning runs of DICESC and the resulting abatement policies. Section five summarizes the findings.

2. The Model

To represent act-then learn, we use a stochastic control version of DICE 2013 [Nordhaus and Sztorc, 2013], DICESC, introduced in Chang and Rutherford [2017]. DICE2013 is a recent version of the original climate IAM authored by Nordhaus [Nordhaus, 1994] and calibrated to be consistent with the IPCC's Fifth Assessment Report [Stocker et al., 2013]. This section presents modifications made to DICE 2013.

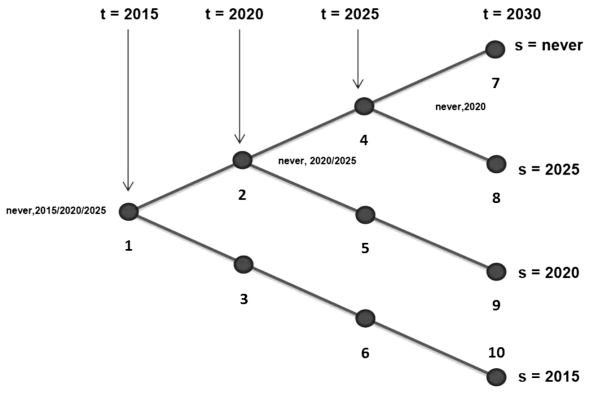

Figure 1. Stochastic Control Recourse Tree

2.1 Recourse: Passive Learning

While maintaining the simple nonlinear program format of DICE, we introduce a recourse tree that integrates the various climate outcomes. The recourse structure represents *states of the world (SoW)*, s, or scenarios associated with the year in which threshold damages take place. *Figure 1.* shows a simple recourse tree made up of four states of the world. For instance, $s = 2015$ represents the unlikely scenario that an irreversible climate catastrophe occurs at year 2015, which the planner only realizes in the subsequent period, 2020. We note that the time horizon of economic activity, t, 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.

The *never* state represents a world in which tipping does *not* take place and corresponds to the recourse path $[1 \rightarrow 2 \rightarrow 4 \rightarrow 7]$. As for SoWs 2020 and 2025, the corresponding paths initially share the nodes on the *never* path but deviate down their respective branches at time 2025 and 2030 when climate tipping is learned. As soon as the economy embarks on a path that experiences climate tipping, a term that represents irreversible catastrophic damages is permanently added to the original damage function in DICE 2013. This term takes the form of a cubic function of atmospheric temperature, relative to the realized tipping threshold.

The recourse structure in DICESC represents two important aspects of the decision making process. First is the *nonanticipative* nature of abrupt climate outcomes modeled. Even in the adverse states of the world, the decision maker is not able to anticipate when, or at which temperature point, tipping takes place. This is because the amount of information (and uncertainty) available in each SoW before the climate tips is identical to that of the *never* SoW. Only when tipping is realized can state specific policies attend to catastrophic damages; ex-ante realization, abatement is largely precautionary.

Realization on the other hand, indicates learning. At every period, a binary event process is realized; i.e. whether or not tipping has occurred. This exogenous arrival of information represents an economy engaging in *passive learning*. However, due to the model's underlying assumption that the economy is capable of recognizing a catastrophic regime shift a period (five years) after the shift takes place, we must assume that such passive learning is really a product of observational learning and *purchased learning* (in the form of R&D) *

Stochastic Control: Active Learning

In a stochastic programming (SP) framework, the hazard rates, and thus the scenario probabilities, are exogenously determined. While the economy could apply periodic updates to the tipping point distribution as it progresses down the *never* path, stochastic programming does not allow for the hazard rate (the likelihood of tipping) to be updated accordingly, so learning remains passive. Following Kelly and Kolstad [1999], Chang and Rutherford [2017] adopts Bayesian learning regarding the location of the threshold to endogenize the likelihood of tipping.

The model adopts a simple prior for single tipping point $\bar{T}^{AT}$, and assumes that the planner is fully aware of the distribution. Note that the maximum temperature threshold in the domain, $\bar{T}_{max}$, is the temperature at which the decision maker believes with certainty that climate tipping occurs.

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

*There is no doubt that observational learning has steadily progressed through the advancement of remote sensing technologies and data trend analyses (especially going from AR4 to AR5), but such assumptions of heightened purchased learning is required to assume the economy realizes tipping in the subsequent period.

Learning on the never path sets the risk domain at time t to be:

$$\bar{T}^{AT} \in [T_t^{AT}(never), \bar{T}_{max}]$$

Given an initial prior distribution, the decision maker engages in *active learning* with respect to the Bayesian prior to compute a period specific hazard rate (hr_t) down the *never* path.

$$hr_t(x) = \frac{\Phi_{\bar{T}_{AT}}[T_{t+1}^{AT}(x)] - \Phi_{\bar{T}_{AT}}[T_t^{AT}(x)]}{\Phi_{\bar{T}_{AT}}[\bar{T}_{max}] - \Phi_{\bar{T}_{AT}}[T_t^{AT}(x)]}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the known Bayesian prior and x, the emissions abatement trajectory over time. The endogeneity of the hazard rate make this optimization problem one of *stochastic control*. Note that this endogeneity stems from global temperature's dependence on abatement trajectory, x. A highlevel problem formulation for stochastic control can be formulated as the following optimization problem:

$$\begin{aligned}
&\max_x \sum_s \Pi_s(x) \cdot F_s(x;\theta_s) \;\; s.t. \\
&\Pi_s(x) = sr_t(x) \cdot hr_t(x), \quad t = s \\
&hr_t(x) = \frac{\Phi_{\bar{T}_{AT}}[T_{t+1}^{AT}(x)] - \Phi_{\bar{T}_{AT}}[T_t^{AT}(x)]}{\Phi_{\bar{T}_{AT}}[\bar{T}_{max}] - \Phi_{\bar{T}_{AT}}[T_t^{AT}(x)]} \\
&sr_t(x) = \prod_{time=0}^{t} (1 - hr_{time}(x))
\end{aligned}$$

where $\Pi_s(x)$ denotes the probability assigned to state of the world (s), which is comprised of the survival rate (sr) and hazard rate at time t. $F_s(x)$ is the climate-economy model for each state of the world, defined by state parameter, θ_s, which is in our case, the time period at which the climate tips. The first order condition of the above stochastic control problem stated in (1), displays the two incentives emissions abatement bears—a damage minimizing

incentive to reduce adverse climate impacts caused by global warming; and a precautionary incentive to *optimize* the likelihood of tipping every period.

$$\underbrace{\sum_{s} \Pi_s(x) \frac{\partial F_s(x)}{\partial x}}_{\text{damage optimizing incentive}} + \underbrace{F_{never}(x) \frac{\partial \Pi_{never}(x)}{\partial x} + \sum_{s^{-never}} F_s \frac{\partial \Pi_s(x)}{\partial x}}_{\text{precautionary incentive}} = 0 \tag{1}$$

Learning and Optimal Abatement

The above first order condition exposes two abatement incentives. The first incentive corresponds to a *damage optimizing incentive*, the incentive that drives abatement when the economy is capable of *passive learning* in the tipping event process. Recall that passive learning gives rise to recourse. Because the optimal policy is an expected abatement trajectory across all states of the world, this policy must optimize over all damage experiences of state-specific decision makers.

The second incentive corresponds to a precautionary incentive, one that drives abatement to change the weights allocated to each state of the world. This incentive comes into being through *active learning* of the Bayesian prior. A first-order decomposition of near-term optimal abatement into the aforementioned incentives can be illustrated using stochastic control and stochastic programming (*Figure 2.*). Stochastic control, which involves both types of learning provides an optimal policy that is significantly more demanding compared to that of DICE 2013. On the other hand, the stochastic programming formulation with a fixed 5% hazard rate every period only exhibits a damage optimizing incentive and outputs a policy trajectory below the stochastic control policy.

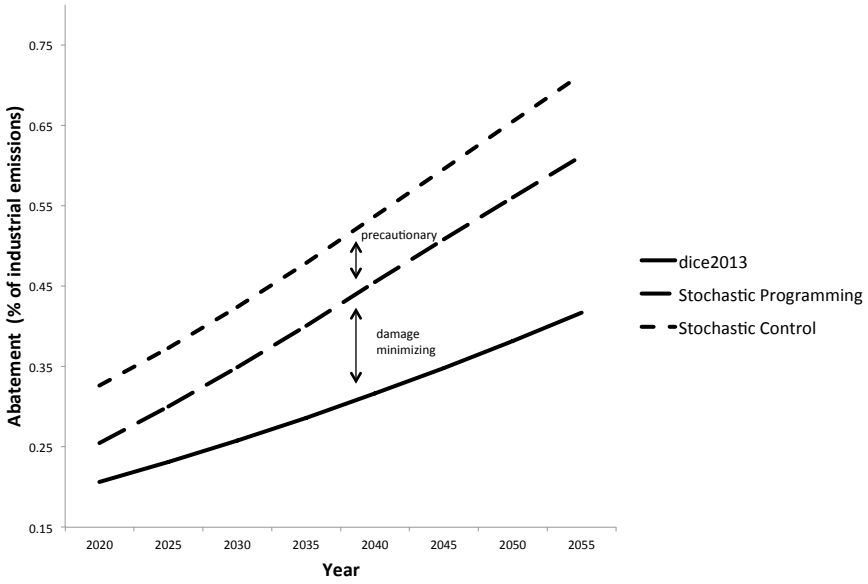

Figure 2. Optimal Near-Term Abatement under DICESC

3. Lack-of-Learning Model Runs

We have thus far determined the optimal policy under stochastic control. What happens then when only a single type of learning is adopted? This section runs two "lack-of-learning" formulations and compares the resulting optimal policy with that of stochastic control.

When only passive learning is employed (DICE-PASSIVE), the formulation is equivalent to a stochastic programming framework with fixed hazard rates. The hazard rates are calibrated to be consistent with the atmospheric temperature trajectory along the business-as-usual (BAU) scenario of DICE 2013. The resulting hazard rates throughout the stochastic time horizon is presented in *Figure 3.* It is no surprise that the resulting hazard rate trajectory for DICE-PASSIVE is much steeper than that of DICESC, reaching almost 100% towards the end of the stochastic time horizon.

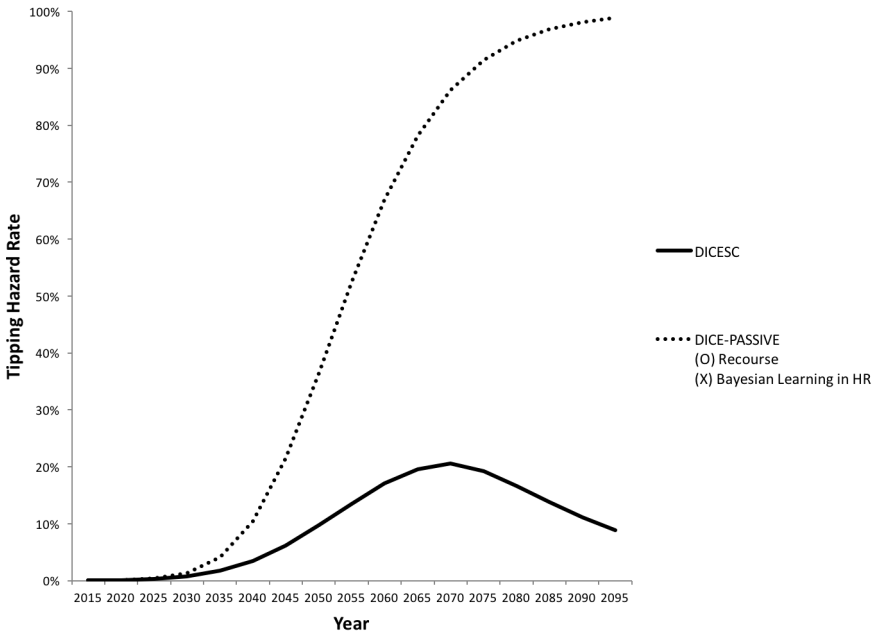

Figure 3. Hazard Rate Calibration to DICE BAU Scenario:

(O) denotes presence of feature; (X) denotes absence of feature

For a comparative study, we perform a similar calibration for the case in which only active learning is adopted to DICE 2013 (DICE-ACTIVE). Since recourse does not exist in this formulation, there exists only the *never* state of the world. Catastrophic damages from climate tipping are thus computed as *expected damages* (as opposed to *realized* damages), allowing the decision maker in each period to hedge between standard DICE 2013 damages (non-catastrophic damages) and a single-period, present value of discounted future threshold damages.

For each period, we determine the present value of tipping damages (PV-Tip) using the recourse tree from DICESC, once again calibrated to be consistent with the temperature trajectory of the BAU scenario in DICE 2013. We also compute the present value of damages down the never state of the world every period (PV-Never). Single-period catastrophic damages are then determined to be the standard DICE 2013 damages multiplied by the ratio of PV-Tip over PV-Never. In other words, single-period catastrophic damages capture the magnitude

of the increase in future damages due to climate tipping, relative to future damages in the absence of tipping, all in present value terms. The resulting damages for catastrophic and non-catastrophic damages in the DICE-ACTIVE framework are presented in *Figure 4.* Given the damage formulation, a decision maker engaged in active learning of the Bayesian prior is able to hedge against catastrophic damages by optimizing the per-period hazard rate.

4. Results and Discussion

We compare the near-term policy output of the three learning formulations; DICESC, DICE-PASSIVE and DICE-ACTIVE. The results are presented in *Figure 5.* One point worth noting is that both lack-of-learning formulations assume potential climate damages that are considerably more severe than damages experienced in the optimal scenario of DICESC. The resulting abatement trajectories nevertheless lie below the DICESC trajectory.

Comparing DICESC with DICE-PASSIVE, we find that the optimal policy is significantly less incentivized in the immediate future in the absence of active learning. This discrepancy, in the immediate future, is driven by the lack of precautionary incentives in the DICE-PASSIVE abatement policy. As the marginal influence of abatement to the state likelihoods decrease over time however, the precautionary incentives in the DICESC policy decreases. In other words, in time, the degrees of freedom in altering the state probabilities will decrease and the role of precautionary abatement will fade. As both policies are increasingly driven by damage optimizing incentives, the discrepancy between the two policies is reduced. Although not shown in the plot, abatement levels of DICE-PASSIVE eventually exceed that of DICESC in the year 2055 due to high hazard rates that follow from the DICE 2013 BAU calibration.

Comparing DICESC with DICE-ACTIVE exposes the role passive learning plays in determining near-term abatement. The initial lack-of-learning policy is not too far off from DICESC's policy. Both policies have a strong incentive to precautionary abate, given how

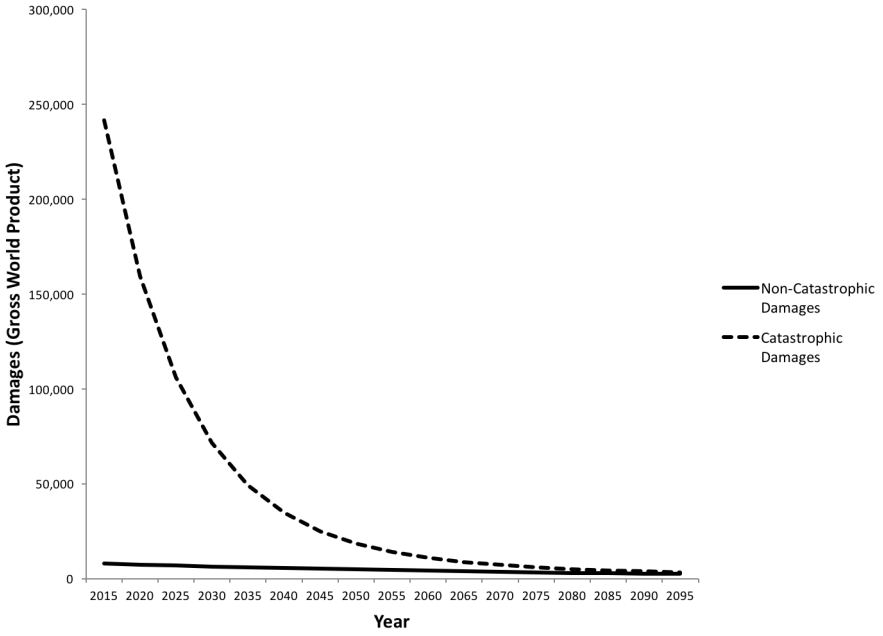

Figure 4. Catastrophic and Non-catastrophic Damages in the DICE-ACTIVE Formulation

Units are in Gross world Product (trillions 2005 USD)

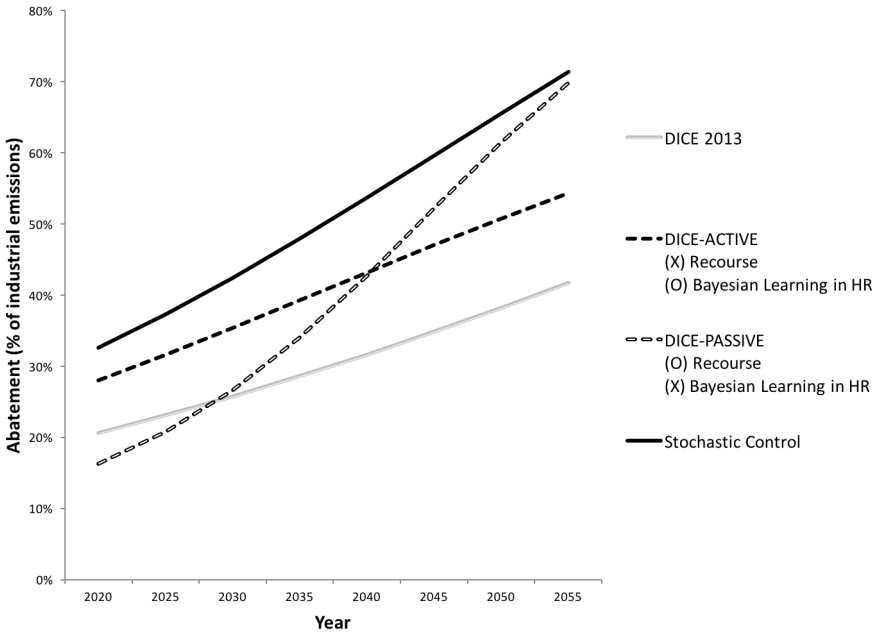

Figure 5. Optimal Near-term Abatement Across Learning Frameworks

(O) denotes presence of feature; (X) denotes absence of feature

severe the damages of immediate climate tipping can be. The discrepancy between the two policies however increases in this case specifically due to the absence of recourse in DICE-ACTIVE. Note that in DICESC, the decision maker adds to the potential tipping states realized as we progress in time. This also means that the severity of tipping damages that are not fully realized in DICE-ACTIVE increase by one every period.

Webster [2002] provides a theoretical model that elucidates how the absence of passive learning leads to lower abatement levels under the risk and uncertainty of climate tipping. His framework consists of a two period model in which decisions today are made under uncertainty, uncertainty which is resolved in the following period. Webster argues that when decisions today change the marginal cost of decisions tomorrow, then the fact that we learn tomorrow influences decisions today in the direction that reduces the *dominant regret* over today's choice, given the expected outcome of learning.

In the context of climate tipping, this means that if more emissions today can lead to catastrophic marginal damages tomorrow through a high-impact, low-probability threshold effect, the dominant regret will stem from learning that the climate has tipped, since staying on the *never* path is already the expectation in the near-term. As a result, when the decision maker commits to sequential observational learning, the dominant regret of having emitted too much carbon will heighten near-term abatement. In other words, the very presence of learning in the model will create additional precautionary incentives to abate in the near-term.

It is clear that neglecting either type of learning can lead to an optimal policy that is severely misleading. Models with no scope of learning will suffer from a lack of incentive to abate in the near-future. A good example of such a model is the MCF method used in Keller et al. [2004]. Keller prescribes a single optimal policy that optimizes welfare throughout

Monte Carlo tipping simulations of the DICE model, and concludes that neglecting uncertainty of tipping points results in a steeper abatement trajectory than when such uncertainty is considered. Since however, there is no scope for learning—neither in the event process of whether or not tipping has taken place, nor in the bayesian prior used to assess the probability of tipping—the resulting optimal policy is insensible to regret over past decisions or the marginal benefits of precautionary abatement. With no scope for learning, there may remain room for improvement in Keller's assessment of his uncertainty effect *.

5. Conclusions

In this paper, we discuss the importance of "act-then learn" strategies in determining the near-term optimal abatement policy. We find that passive learning in the history of the event process creates abatement incentives in the near-term to reduce the "regret" from learning that the climate has tipped in the subsequent period. We also find that active learning in the Bayesian tipping prior creates a strong precautionary abatement incentive in the immediate future. Although the Monte Carlo Fitting method does not suffer from the improper handling of uncertainty *per se*, the method nevertheless suffers from a lack of abatement incentives due to the absence of learning. Decision making under uncertainty on its own may not suffice in prescribing a near-term optimal carbon abatement policy, especially when observational learning (passive learning) functions as a large part of the sequential learning process.

*We note that Keller et al. [2004] considers the risk of thermohaline collapse. We are cautious in making this claim due to the long transition time scale of 100 years, as assessed by Lenton et al. [2008], which puts the relevance of passive learning to question. We also note that in modeling the tipping of the West Antarctic ice sheet, Diaz and Keller [2016] employs MCF in the stochastic control framework introduced by Chang and Rutherford [2017]

Chapter 4.

Managing the Risk of Climate Thresholds: Optimal Policies Across Models of Ambiguity Aversion

1. Introduction

One component of climate impacts that impedes consensus is the risk and uncertainty of climate catastrophes. The current literature on climate tipping —catastrophic climate outcomes triggered by atmospheric temperatures exceeding a threshold—indicate that neglecting such outcomes can present severely misleading optimal policies [Pindyck, 2013, Weitzman, 2011]. Accordingly, policy relevant tipping elements such as the Atlantic Thermohaline Circulation (THC) collapse, West Antarctic ice sheet melting (WAIS) and Boreal forest dieback have entered numerous climate integrated assessment models (IAMs) [Lenton et al., 2008]. Although models are constructed using available data and the best of expert knowledge, the quantification of tipping impacts is characterized by large amounts of uncertainties, which not only consist of scientific uncertainty that stems from the complex climate system, but also socioeconomic uncertainty, how the world would respond to an adverse regime shift in climate [Heal and Millner, 2014].

Despite the slow rate at which climate uncertainty is resolved, the climate science literature has experienced a rapid increase in the number of "improved" scientific models available. But as a result, expert knowledge and model specific findings on climate impacts, let alone

how the climate responds to increased greenhouse gas concentrations, are often found to be inconsistent with one another.

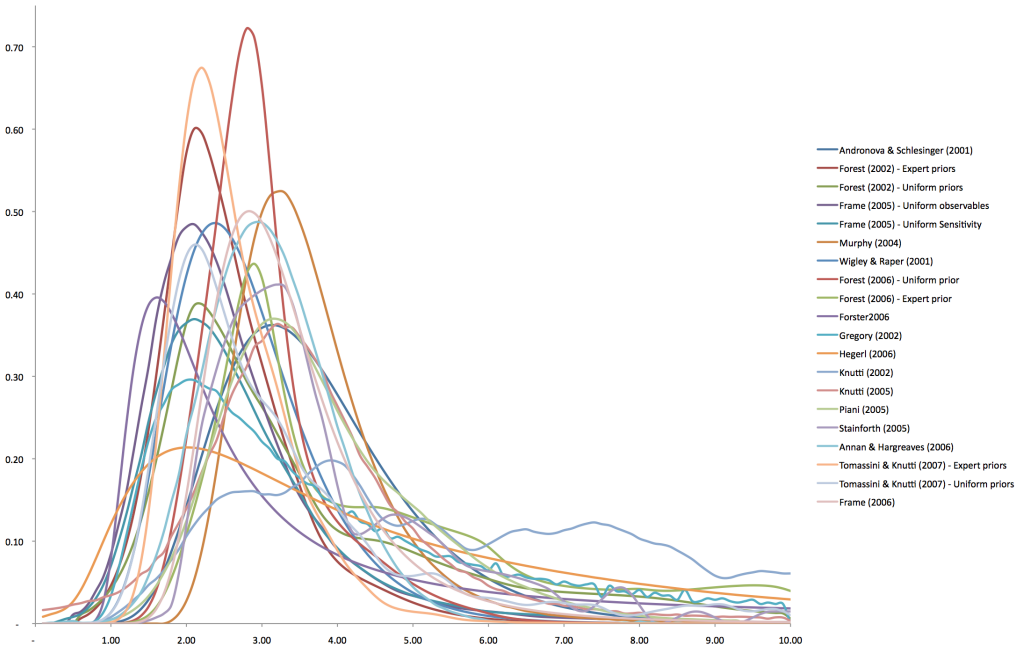

Figure 1. Estimates of the probability distribution for climate sensitivity from Millner et al. [2013]

In other words, the uncertainty literature in climate change has evolved enough to separate the notion of risk from uncertainty. The term *risk* is used to characterize environments in which the probability distribution of a random event taking place is known with certainty. In contrast, the notion of *uncertainty* (also known as Knightian uncertainty or ambiguity) characterizes an instance in which the decision maker is unsure of the probability distribution that describes the likelihood of the risky event taking place.

For instance, Millner et al. [2013] feature probability density functions for climate sensitivity from a variety of climate studies, originally collated by Meinshausen et al. [2009] (see *Figure 1*). Each probability density (model) obtained is based on the study's unique representation of the climate's physical and biological processes. The type of data used to estimate climate sensitivity also differs among studies. For this reason the decision maker (DM) can-

not objectively rank models to settle order of validity, nor establish dependency of one model to another, framing the DM's welfare maximization problem as one of *model ambiguity* [Heal and Millner, 2014].

Decision making under *ambiguity aversion* was first introduced in Ellsberg [1961], to point out that real world decision making when faced with uncertain beliefs cannot be represented by the axiomatic foundations of expected utility theory, namely the Savage axioms [Savage Leonard, 1954]. Leonard Savage had argued (with caution) that when preferences obey his axioms, decision making under uncertainty could be represented under a generalized expected utility framework in which agents would form *subjective probabilities*, even in the total absence information about the likelihoods of an outcome taking place. Daniel Ellsberg's well known *Ellsberg paradox* however points out that people in real life are averse to uncertainty, and would prefer lotteries with established probabilities over ones with no information on the event likelihoods, thereby violating Savage's axioms that characterize rational choice.

The general literature on ambiguity aversion as a result [Al-Najjar and Weinstein, 2009, Machina and Siniscalchi, 2014, Hansen and Sargent, 2008], presumes that choices observed in the Ellsberg paradox are rational [Al-Najjar and Weinstein, 2009]. The recent ambiguity literature in climate economics establishes the relevance of ambiguity aversion in the various elements of (mainly scientific) climate uncertainty [Lange and Treich, 2008, Heal and Millner, 2014], attends to the challenges of modeling uncertainty and ambiguity aversion [Gilboa and Schmeidler, 1989, Klibanoff et al., 2009, Schmeidler, 1989, Traeger, 2011, Millner et al., 2013], and lastly provides assessments of optimal policy that deal with model uncertainty using both analytical and numerical models [Drouet et al., 2015, Berger et al., 2016, Lemoine and Traeger, 2016a, Rudik, 2016, Millner et al., 2013].

This paper contributes to two aspects of the ambiguity literature in climate economics.

First, I present how ambiguity in the likelihood of climate tipping can be modeled based on four different definitions of ambiguity aversion in a general stochastic control setting with *multiple prior distributions*. Drouet et al. [2015] finds that among popular sources of uncertainty modeled in the climate literature (impacts, mitigation costs and model uncertainty), uncertainty in damages has critical impacts on expected gross world product, especially in a high emissions scenario over the twenty first century. Secondly, I add to the literature of assessing optimal policy under model uncertainty. Using a stochastic control version of Nordhaus' DICE model [Nordhaus, 1994, Nordhaus and Sztorc, 2013] introduced in Chang and Rutherford [2017], I provide numerical assessments of the optimal abatement policy under each ambiguity averse attitude. The paper's modeling of catastrophic damages via climate tipping lastly relates to the literature on the risk of tipping points and near-term optimal policies [Chang and Rutherford, 2017, Diaz and Keller, 2016, Lemoine and Traeger, 2016b, 2014, Keller et al., 2004, Lontzek et al., 2012, Cai et al., 2012, Lenton, 2012].

For all definitions of ambiguity aversion considered in this paper, I find robust precautionary incentives to heighten optimal policy throughout the twenty first century. Furthermore, in a multi-prior, stochastic control setting with two type of priors—the per period tipping hazard rate and the probabilities assigned to states of the world—I find that hedging strategies can differ significantly depending on which prior aversion is defined on. I note that, with respect to the approach used to analyze the effect of model uncertainty to policy, the nature of my research bears closest resemblance to the work of Berger et al. [2016], which considers uncertainty in the likelihood of climate tipping in a stochastic control framework. Berger et al. uses a derivative model of *smooth ambiguity preference* introduced by Klibanoff et al. [2005] in a multiple-prior setting. Millner et al. [2013] and Lemoine and Traeger [2016a] also use models based on generalizations of smooth ambiguity aversion.

When the aversion model allows for a single Bayesian prior to be assumed per period, as in the Maxmin [Gilboa and Schmeidler, 1989] and Choquet ambiguity [Schmeidler, 1989] aversion models, I find that transition points between priors can impose threshold effects on policy even before tipping occurs, especially when the agent is also *risk averse*. This is because an agent averse to both uncertainty and risk exhibits strong incentives to delay reaching temperature points at which the assumed Bayesian prior takes an abrupt turn for a "worse" prior. This paper is organized as follows: the next two sections respectively present models of ambiguity and risk aversion. Section four provides an assessment of optimal policy across ambiguity attitudes. In section five, I devise optimal policies under aversion to both risk and ambiguity. The last section provides discussion on results.

2. Models and Definitions of Ambiguity Aversion

The setting in this paper is based on the stochastic control version of DICE, DICESC, introduced in Chang and Rutherford [2017]. Chang and Rutherford consider an endogenous hazard rate of climate tipping under a known Bayesian prior. I incorporate model uncertainty to this hazard rate in a multiple prior setting. In Chang and Rutherford [2017], we follow Stocker et al. [2013], Lenton et al. [2008] to assume at least one tipping element is triggered with certainty when global temperatures exceed the expected equilibrium climate sensitivity of 3 °C, and model for a single tipping element with an unknown temperature threshold. This assumption is also consistent with the work of Kriegler et al. [2009], which obtains lower bounds to the probability of triggering at least one large-scale tipping point, including the Greenland and West Antarctic ice sheets, Atlantic meridional overturning circulation (AMOC) and the Amazon rainforest. Among policy relevant tipping elements, *West Antarctic ice-sheet melting (WAIS)* is an adequate instance to model using DICESC due to the its short transition time and associated uncertainty domain of the tipping threshold.

I consider three prior distributions of risk to construct a multiple prior setting. The baseline prior is a normal distribution calibrated so that the upper and lower tail decline to $\epsilon > 0$ at each end of the domain, 0.83 °C and 3.0 °C, respectively. Note that we define uncertain threshold, $\bar{T}^{AT}$, as the increase in global atmospheric temperature since year 1900 and set the domain of the prior distribution to be $\bar{T}^{AT} \in [\bar{T}_0, \bar{T}_{max}] = [0.83, 3.0]$ °C. Here, $\bar{T}_0$ corresponds to the initial increase in atmospheric temperature since 1900, which is parameterized as 0.83°C in DICE, and $\bar{T}_{max}$, the expected value of climate sensitivity, for which we assume tipping takes place with certainty. Also note that, in the baseline distribution a tipping threshold of 1.915 °C is assigned the highest likelihood, with equal weights assigned to the worst and best case scenarios. Following Lemoine and Traeger [2014], Kelly and Kolstad [1999], the hazard rate under each prior is defined as:

$$hr_t(x) = \frac{\Phi_{\bar{T}_{AT}}[T^{AT}_{t+1}(x)] - \Phi_{\bar{T}_{AT}}[T^{AT}_{t}(x)]}{\Phi_{\bar{T}_{AT}}[\bar{T}_{max}] - \Phi_{\bar{T}_{AT}}[T^{AT}_{t}(x)]} \tag{1}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the known Bayesian prior and x, the emissions abatement trajectory over time. Based on the hazard function, the probability of tipping in the current period is the normalized probability mass that corresponds to the temperature rise experienced in the current period, under a specific prior distribution.

To the baseline, I apply the notion of *increasing risk* defined by Rothschild and Stiglitz [1970] and add two mean-preserving spreads (MPS). For the first MPS, I consider a uniform prior over the baseline domain. For the second, I keep the first and second moments fixed to those of the uniform prior and construct a one-sided truncated normal prior (on temperature domain $[0.83, \infty]$). *Figure 2* plots the three priors used in the current analysis.

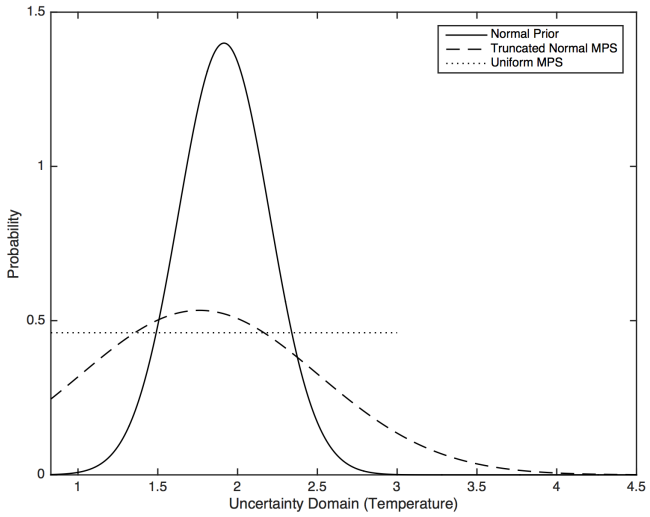

Figure 2. Tipping Prior Distributions

Given $hr_t(x)$, the conditional probability of climate tipping during period t (assuming that there has been no tipping to that point in time), is used to assign probability Π_s to each state of world, s; i.e. $\Pi_s(x) = sr_t(x) \cdot hr_t(x), \;\; t = s$, where $sr(t)$ is the survival rate, the probability that the climate does not tip until time $t = s$. Note that the states of the world are defined by *when* tipping takes place within the economic time horizon. The stochastic control problem given known Bayesian prior, p, is thus defined as:

$$\max_x \sum_s \Pi_s^p(x) \cdot F_s(x;\theta_s) \;\; s.t. \tag{2}$$

$$\Pi_s^p(x) = sr_t^p(x) \cdot hr_t^p(x), \quad t = s \tag{3}$$

$$hr_t^p(x) = \frac{\Phi^p_{\bar{T}_{AT}}[T^{AT}_{t+1}(x)] - \Phi^p_{\bar{T}_{AT}}[T^{AT}_t(x)]}{\Phi^p_{\bar{T}_{AT}}[\bar{T}_{max}] - \Phi^p_{\bar{T}_{AT}}[T^{AT}_t(x)]} \tag{4}$$

$$sr_t^p(x) = \prod_{time=0}^{t} (1 - hr^p_{time}(x))$$

where F_s represents the climate-economy model over state s, defined by state parameter, θ_s, which is in our case, the time period at which the climate tips. Note that endogeneity of

the state probabilities $\Pi_s^p(x)$ is based on current atmospheric temperature's dependence on previous period abatement levels.

In this section, I present four aversion attitudes toward ambiguity in a stochastic control framework; namely, Maxmin [Gilboa and Schmeidler, 1989], Choquet [Schmeidler, 1989], Smooth [Klibanoff et al., 2005] and Choquet-Smooth ambiguity aversion. The optimization problem stated in $(2) - (4)$ is solved for each model, however we no longer assume perfect information in the Bayesian prior and deal with a multiple prior stochastic control problem. Ultimately, the different aversion attitudes toward ambiguity boils down to how we deal with the set of multiple priors, to determine the hazard rate (and hence the state probabilities) used in our stochastic control problem.

2.1 Maxmin Ambiguity Aversion

The idea of optimal policies under ambiguous distributions appeared in early attempts by Scarf et al. [1958], Gilboa and Schmeidler [1989], which led to the development of a field in Operations Research known as *distributionally robust* multiperiod stochastic optimization. Reviewed extensively in Analui and Pflug [2014], this field of study is motivated by the decision making process that consider all model possibilities to output a robust decision strategy. Synonymous names for stochastic problems that deal with multiple priors are, model uncertainty problem, distributionally robust problem, and *maxmin stochastic optimization* [Analui and Pflug, 2014].

The term maxmin stochastic optimization (MSO) is used to characterize a solution to the model ambiguity problem for which the objective is to hedge or maximize welfare against the worst expected outcome, resulting from a set of permissible distributions [Shapiro and Kleywegt, 2002]. In the context of climate policy, a general application of MSO to tipping

uncertainty models can be summarized in the following optimization problem:

$$\max_{x \in X} \left[\min_{p \in \mathcal{P}} \mathbb{E}_p \big[F(x) \big] \right] \equiv \max_{x \in X} \left[\min_{p \in \mathcal{P}} \sum_s p_s(x) \cdot F_s(x) \right] \tag{5}$$

where $\mathcal{P}$ denotes the set of probability distributions p, over states $s \in S$, and $F(x)$, the model as a function of emissions abatement level x. In essence, the solution concept of maxmin takes into account all priors that are consistent with historical data, only to maximize for the least favorable prior. Again, for problems in which a systematic ranking of prior estimates is nearly impossible—as in the estimation of the climate sensitivity probability distribution *Figure 1*—every prior is given equal status. In this context, focusing on the least favorable prior (estimate or combinations of priors) reflects extreme ambiguity aversion [Heal and Millner, 2013]. The resulting ambiguity aversion model is called Maxmin Expected Utility (MEU), or sometimes the Multiple-Priors (MP) model. I formally characterize the closed, convex and *rectangular* nature of the prior set used in the other aversion attitudes, using the MEU model as an example.

Applying MEU to uncertain hazard rates involves constructing a *least favorable prior* as a combination of *one-step-ahead* priors—primitive priors displayed in *Figure 2*, i.e. $\{\Delta_{uniform}(S), \Delta_{tnormal}(S), \Delta_{normal}(S)\} \in \mathcal{P}$—assumed every period. p_s, in formulation (5), is then the likelihood of state s computed using a combination of probability measures which in turn, are based on a set of bayesian priors over the temperature domain.

There consequently exist two kinds of priors; one on the location of the tipping point (denote $\mathcal{P}^{TATM}$), and another on the probability distribution over states of the world (denote $\mathcal{P}^S$). Note that $\mathcal{P}^S$ is a function of both $\mathcal{P}^{TATM}$ and abatement levels x. To formally characterize the set of priors $\mathcal{P}^S$, I use Epstein and Schneider [2003]'s model of dynamic choice by a ambiguity averse decision maker. Define $\mathcal{P}^S$ to be a $\mathcal{F}_t$*-rectangular set* on $(\Omega, \mathcal{F}_T)$, given state

space $\Omega = \{never, 2015, 2020, \ldots, T\}$ and information structure represented by the filtration $\{\mathcal{F}_t\}_t^T$ where, $\mathcal{F}_{2015} \subset \ldots \subset \mathcal{F}_T$. Note that $\mathcal{F}_t$ is generated by a finite partition; that is, $\mathcal{F}_t(s)$ denotes the partition component containing $s \in \Omega$. A set of priors $\mathcal{P}^S : \Omega \to \Delta(\Omega, \mathcal{F}_T)$, where $\Delta(\Omega, \mathcal{F}_T)$ denotes the set of probability measures on the $\sigma-$algebra $\mathcal{F}_t$, is *rectangular* if for $\forall t, s$

$$\mathcal{P}_t^S(s) = \left\{ \int_\Omega p_{t+1}(s^{'})dm : p_{t+1}(s^{'}) \in \mathcal{P}_{t+1}^S(s^{'}) \ \ \forall s^{'} \in \Omega, m \in \mathcal{P}_t^{+1}(s) \right\}, \tag{6}$$

$$for \quad \mathcal{P}_{t+1}^S(s) = \{p_{t+1}(s) : p \in \mathcal{P}^S\} \tag{7}$$

$$\mathcal{P}_t^{+1}(s) = \{p_t^{+1}\big(s, p^{TATM}(x)\big) : p^{TATM} \in \mathcal{P}^{TATM}, x \in [0,1]\} \tag{8}$$

where for any measure p on $(\Omega, \mathcal{F}_T)$, $p_t(s) = p(\cdot|\mathcal{F}_t)(s)$ is its $\mathcal{F}_t$-conditional and p_t^{+1} is the restriction of p_t to filtration at time t, $\mathcal{F}_t$.

We can think of $\mathcal{P}_t^{+1}$ as the set of conditional *one-step-ahead* probability measures describing the likelihood of a hazard in the next period. This set consists of probability measures assigned to each binary partition component—the hazard rate and survival rate—which is a function of not only the type of distribution over the temperature domain (p^{TATM}) used to characterize the risk of tipping, but also the increase in global atmospheric temperature, which is in turn a function of emissions abatement x at any time period. As a result, even with a finite number priors in set $\mathcal{P}^{TATM}$, there exists an infinite number of one-step-ahead probability measures in the set $\mathcal{P}_t^{+1}(s)$ for $\forall(t, s)$.

Applying Maxmin stochastic optimization in this context corresponds to determining (p_t^{TATM}, x_t) to assign the greatest one-step-ahead probability measure to the *least favorable outcome* every time period. If we focus on the partial equilibrium problem of choosing a temperature based prior $p_t^{TATM} \in \mathcal{P}^{TATM}$, in essence, the decision maker chooses the prior

that assigns the greatest hazard rate given the amount of temperature increase the economy experiences between time periods $[t, t+1]$. In the context of the full optimization scheme (5) in which x is a control variable, the decision maker chooses abatement trajectory, $\{x\}_{t=2015}^{T}$ to maximize expected utility while being mindful of the ambiguity averse process of selecting one-step-ahead priors each period.

Epstein and Schneider [2003] show that the rectangular set of priors, $\mathcal{P}^S$, can be uniquely constructed using a primitive set of one-step-ahead correspondences $P_t^{+1} : \Omega \to \Delta(\Omega, \mathcal{F}_{t+1})$. That is, by setting $\mathcal{P}_t^{+1} = P_t^{+1}$ for $\forall t, s$, and using backward induction under relation:

$$\mathcal{P}_t^s(s) = \left\{ \int_\Omega p_{t+1}(s^{'})dm : p_{t+1}(s^{'}) \in \mathcal{P}_{t+1}(s^{'}) \ \ \forall s^{'} \in \Omega, m \in P_t^{+1}(s) \right\} \tag{9}$$

rectangular set, $\mathcal{P}^s$, is uniquely determined to be consistent with the set of one-step-ahead conditionals, P_t^{+1}, $\forall t$; i.e.

$$\mathcal{P}^S = \left\{ p \in \Delta(\Omega, \mathcal{F}_T) : p_t^{+1} \in P_t^{+1} \ \ \forall t \text{ and } \forall s \right\} \tag{10}$$

for which the set of one-step-ahead conditionals, P_t^{+1}, is in turn determined by tipping prior, $p^{TATM} \in \mathcal{P}^{TATM}$, and abatement level, x:

$$P_t^{+1} = \{ p_t^{+1}(x, p^{TATM})) : p^{TATM} \in \mathcal{P}^{TATM}, x \in [0,1] \} \tag{11}$$

Although rectangularity is a notion of recursive preferences, it helps formally characterize the set of priors that are available to the agent in a multiple prior stochastic control setting (stochastic control and stochastic programming employs forward optimization). In DICESC, rectangularity is trivially satisfied; it is clear that the dynamic set of priors is rectangular with respect to the information structure. The composition of priors and one-step-ahead condi-

tionals with respect to the binary partition structure at every period, does not conflict with the decision maker's prior beliefs; that is, given the set of priors $C \subseteq \mathcal{P}^S$ that satisfy $p_s(\cdot) = 0$ for $\forall s < t$ or equivalently, $\mathbb{P}(\bar{t} > T_{AT}(t)) = 1$ for uncertain tipping point $\bar{t}$, the composition of priors $P(E) \cdot Q(\cdot|E), \ \forall P, Q \in C$ defined by relation (9) is also in C. Rectangularity is trivially satisfied by definition *.

Given multiple priors, hazard function in the MEU is formulated as:

$$hr^*(t) \equiv \max_{p \in P^{+1}} \ hr^p(t) \tag{12}$$

$$\equiv \max_{p \in P^{+1}} \ \frac{F^p_{\bar{T}_{AT}}[T_{AT}(t+1)] - F^p_{\bar{T}_{AT}}[T_{AT}(t)]}{F^p_{\bar{T}_{AT}}[\bar{T}_{max}] - F^p_{\bar{T}_{AT}}[T_{AT}(t)]} \quad \forall t \in T \tag{13}$$

for which the set of next-step primitive correspondence simply encompasses the three priors displayed in *Figure 2*. We can thus restate the model objective as:

$$\max_{x \in X} \ \sum_s \Pi^*_s(x) \cdot F_s(x) \tag{14}$$

$$\text{where } \ \Pi^*_s(x) \equiv sr^*(t) \cdot hr^*(t) \tag{15}$$

Lastly, a note on constructing the asserted set $\mathcal{P}^S$; while recursive preferences use backward induction based on relation $(9)^\dagger$, DICESC's construction employs a forward induction procedure. At every time period of the economic time horizon, the planner must assess the likelihood of tipping using all primitive priors, which creates an additional dimension to the decision process—the decision of how much to abate, or equivalently, how much temperature increase the economy allows between periods, determines not only the hazard rate conditional on the prior, but also the prior (in $\mathcal{P}^S$) itself. At the lower end of the temperature

*See Al-Najjar and Weinstein [2009] for further details in checking for rectangularity.

$^\dagger \mathcal{P}_T(s)$ becomes a single function $p(\omega) = 1_\omega(s)$ for $\forall s \in \mathcal{F}_T$, as the economy is aware of which event in $\mathcal{F}_T$ was realized at T.

distribution, a uniform distribution is *least favorable* as it assigns most weight to the most damaging scenarios and thus acts as the effective $p \in \mathcal{P}^{+1}$ in the model. At the high end of the distribution however, the hazard function is most likely based on the normal distribution as a unit increment in temperature results in the highest likelihood of tipping.

MEU describes the planner's consideration of the current state of the economy, and the worst likelihoods of a catastrophic outcome pertaining to the prescribed actions every period, specifically under the assumption that there exists *no resolution of distributional ambiguity* throughout the time horizon.

2.2 Choquet Expected Utility

A model of ambiguity aversion that has not been used widely in the climate economics ambiguity literature, is the rank-dependent expected utility, or the *Choquet expected utility model (CEU)*, introduced in Schmeidler [1989]. I provide a summary of the CEU framework based on Machina and Siniscalchi [2014]. Unlike MEU, in which aversion behavior is captured by the choice of prior, aversion to ambiguity in CEU is represented through a *non-additive probability measure*. Such a probability measure assigns *capacities*, $v(\cdot)$, to the subsets of the state space S, while satisfying the normalization $v(\emptyset) = 0$, $v(S) = 1$ and the monotonicity property; i.e. $E \subseteq F$ implies $v(E) \leq v(F)$. As the name of the measure suggests, the sum of the probability over states need not sum to unity. It is through this residual probability mass that ambiguity is expressed in CEU. Machina and Siniscalchi [2014] uses the Ellsberg example to elaborate on this point. For instance, an agent may have the following beliefs when selecting a black ball versus a yellow ball,

$$v(black, yellow) = 2/3, \quad v(black) = v(yellow) = \epsilon \in [0, 1/3).$$

The existence of a residual probability, $v(black, yellow) - v(black) - v(yellow) > 0$, is what represents ambiguity in CEU, as it is indicative of the inability of the agent to distribute the residual mass between the two colors.

To resolve the issue of non-additivity, Schmeidler [1989] introduces the *Choquet integral*. Given function $g : S \to \mathbb{R}$ that maps the state space to an (discrete) outcome space, and a ranking of the outcomes $\alpha_1 > \ldots > \alpha_n$, $g(s)$ is evaluated using Choquet integral:

$$\int g(s)dv \equiv \alpha_1 \cdot v(s : g(s) = \alpha_1) + \sum_{i=2}^{n} \alpha_i \cdot \big[v(s : g(s) \geq \alpha_i) - v(s : g(s) \geq \alpha_{i-1})\big].$$

In a stochastic control setting, assuming that the agent is able to assign capacities to each state of the world, we define the residual mass as being the residual probability after assigning state likelihoods under CEU; i.e. $1 - \sum_s p_s^{CEU}(x)$ in formulation (5). Furthermore, in a multiple prior setting, the agent updates ambiguous beliefs using *full Bayesian updating* [Machina and Siniscalchi, 2014, Fagin and Halpern, 1991], in which the agent considers the *lower envelope* of the dynamically consistent set of priors each period (note again that this set is endogenously constructed as a function of both abatement and primitive priors). That is, capacity $v(s)$ is defined as the $v(s) = \inf\{p^s(s) : p^S \in \mathcal{P}(x)^S\}$. As a result, the set, $\mathcal{P}(x)^S$, is characterized with priors that assign *at least* $v(s)$ to state s, given abatement vector, x. Lastly, the residual probability is distributed to the most averse state of the world via the Choquet integral. Ambiguity aversion in the CEU model is thus defined by the agent's conservative stance on assessing the probability of each tipping state of the world (assuming the *infimum* measure, given ambiguous information), and also by the distribution of the residual probability (assuming the*worst* distribution, given ambiguous information).

In DICESC, this means that priors are assumed—for each time period and each state of the world—to minimize state likelihoods which are made up of the survival rate (sr) and

hazard rate (*hr*):

$$p_s^{CEU} = sr^{CEU}(t|t=s) \cdot hr^{CEU}(t|t=s). \tag{16}$$

To minimize each component of the likelihood function, the agent must assume a combination of the "worst" priors (equivalent to MEU priors) up until the time of tipping to minimize the survival rate, and in contrast assume *the most convenient* prior to minimize the hazard rate at the time of tipping. The residual probability mass, if any, is swept under the likelihood of the state in which tipping occurs in year 2015 (the worst state of the world). In contrast to MEU, in which a unique single prior is used, there are many priors in set $\mathcal{P}^S$ that is in effect in CEU. The resulting probability assignment for each state specific prior p^s, assigned at time 0, do not need to sum to unity; i.e. $\sum_s p_0^s(s) \leq 1$.

Bear in mind that the application of CEU in DICESC deals with the uncertainty in state likelihoods as a result of uncertainty in the tipping hazard rate, as opposed to dealing directly with uncertainty in the hazard rate (which is the case in MEU). The reason for this is that, CEU is equivalent to MEU when applied directly to the hazard rate. Since the filtration partition is binary at every time period (occurrence and absence of tipping), assigning probabilities to each partition that correspond to the lower envelope of the prior set leads the agent to consider two different priors that respectively put the least amount of mass on both partitions. The Choquet integral however, assigns the residual probability mass to the tipping partition, effectively making the resulting prior equivalent to that of the MEU model.

2.3 Smooth Ambiguity Preferences Model

The Smooth Ambiguity Preferences model (SAP), introduced in Klibanoff et al. [2005], characterizes an ambiguity averse agent that is uncertain about which probability measure (one-

step-ahead prior, $p_t^{+1}(s)$), best describes the likelihood of the tipping hazard. In this model, aversion to ambiguity is expressed using a two-stage approach [Machina and Siniscalchi, 2014]. Consider act $f : S \to \chi$, a mapping of states (S) to outcomes (χ). Act f can be thought of as a "policy" and is evaluated in SAP as:

$$W\big(f(\cdot)\big) = \int_{\mathcal{P}^{TATM}} \phi\bigg(\int U\big(f(\cdot)\big) dp^{TATM} \bigg) dM(p^{TATM}) \tag{17}$$

where $U(\cdot)$ is a vNM utility function, $M(\cdot)$, the second-order prior over the primitive priors and μ, the primitive priors over the state space, S. The inner integral represents the expected utility under a specific tipping prior, while the outer integral assigns probabilities or weights based on belief, to the outcomes of each prior. Aversion, or preference to ambiguity is expressed through the shape of the second-order utility function, ϕ—if $\phi(\cdot)$ is concave, the expression in (17) represents aversion to ambiguity, whereas convexity and linearity respectively imply preference and neutrality to ambiguity. Under ambiguity aversion, the agent applies the certainty equivalence operator over the expected utilities corresponding to each primitive prior, putting more importance on the prior assumptions that lead to "averse" expected outcomes.

Millner et al. [2013] defines an "ambiguity-adjusted" distribution, $M(\cdot)$ that assigns weights to each primitive prior to express the agent's aversion to ambiguity. Millner et al. assumes an isoelastic for the second-order weighting function; given coefficient of relative ambiguity aversion, η, function certainty equivalence function is defined as:

$$\phi(U) = \begin{cases} \frac{U^{1-\eta}}{1-\eta}, & \eta < 1 \\ -\frac{(-U)^{1+\eta}}{1+\eta} & \eta > 1 \end{cases} \tag{18}$$

and the "ambiguity-adjusted" second-order probabilities that determine the weights put on each prior, as a function of primitive prior and abatement levels, $p_m(p^{TATM}, x)$, are defined as follows:

$$p_m(p^{TATM}, x) = \frac{\phi'\Big(EU(p^{TATM}, x)\Big)}{\sum\limits_{q \in P^{TATM}} \phi'\Big(EU(q^{TATM}, x)\Big)} = \frac{\phi'\Big(\int U\big(f(\cdot, x)\big) dp^{TATM}\Big)}{\sum\limits_{q \in P^{TATM}} \phi'\Big(\int U\big(f(\cdot, x)\big) dq^{TATM}\Big)} \tag{19}$$

Note that due to the concavity of the second-order weight function, the first-order derivative, $\phi'(\cdot)$ is decreasing and as a result, assigns more weight to priors that generate lower expected utilities, under a fixed level of abatement, x.

Lastly, in contrast to MEU and CEU, for which the set of one-step-ahead priors consisted of only the three primitive priors displayed in *Figure 2*, SAP's one-step-ahead prior set consist of priors characterized by the convex combination of the primitive set.

2.4 Smooth-Choquet Ambiguity Aversion

In CEU, the residual probability mass resulting from the Choquet integral is simply assigned to the worst state of world. In this subsection, I introduce *smooth CEU (SCEU)*, in which the residual probability mass is distributed to all states of the world using *risk-adjusted* weights. The main objective of using SCEU is to assume a hedging attitude less extreme than the one assumed in CEU, and to also involve the notion of risk aversion in forming attitudes toward ambiguity. Keep in mind that when applying CEU to uncertainty in state likelihoods, the distribution of the residual mass is decisive in determining the ambiguity attitude.

Let us define $U(s)$, as the state dependent utility, and $\phi(\cdot)$, as the certainty equivalent

operator introduced in the SAP model:

$$\phi(U(s)) = \begin{cases} \frac{U(s)^{1-\eta}}{1-\eta}, & \eta < 1 \\ -\frac{(-U(s))^{1+\eta}}{1+\eta} & \eta > 1 \end{cases} \tag{20}$$

Then using a similar weighting system as in the SAP model, I can create risk-adjusted weights, $p_r(s,x)$ to be used in distributing the residual probability mass:

$$p_r(s,x) = \frac{\phi'\Big(U(s,x)\Big)}{\sum\limits_{s\in S}\phi'\Big(U(s,x)\Big)} \tag{21}$$

Similar to the SAP model, the SCEU model also employs a two-stage approach. The first stage involves the assignments of CEU state likelihoods before the Choquet integral, p_s^{CEU}, computed as a function of one-step-ahead primitives and abatement levels (equation (16)). The second step is to assign the residual probability mass, $(1-\sum\limits_{S} p_s^{CEU})$ to all states based on risk-adjusted weights. The state likelihoods in SCEU results in the following expression:

$$p_s^{SCEU}(x) = p_s^{CEU}(x) + p_r(s,x)\cdot\big(1-\sum_{S} p_s^{CEU}\big) \tag{22}$$

The SCEU model still exhibits the conservative beliefs of the CEU model agent (assuming the *infimum* probability measure, given ambiguous information), while adding scope for smooth, risk-aversion in the distribution of the residual probability.

3. Results

3.1 Optimal Policy

In *Figure 3*, I display the optimal near-term abatement policy across models of ambiguity aversion. On top of the four attitudes toward ambiguity, I plot the policies for the original DICE model (denoted DICE 2013), the stochastic control version of DICE (DICESC) and the case under the principle of Insufficient Reason; i.e. equal weights for each one-step-ahead prior. The policies output by ambiguity aversion indicate a robust incentive to heighten policy in the near-term.

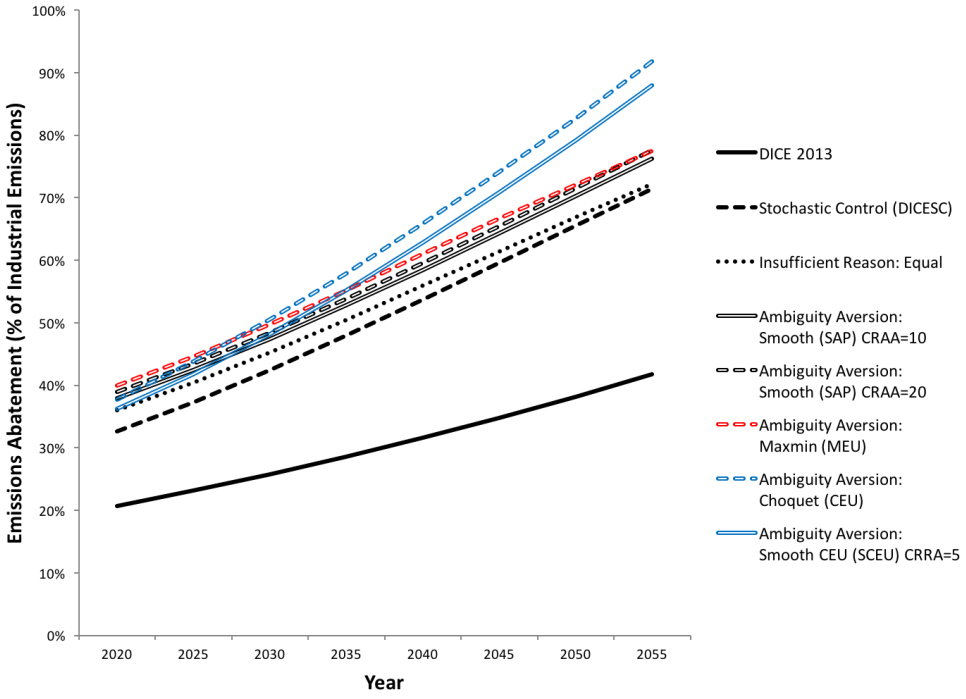

Figure 3. Optimal Near-Term Abatement Across Models of Ambiguity Aversion

There are three points to consider. Firstly, abatement levels are greatest under the CEU and MEU models. Since the MEU model corresponds to a special case of the smooth ambiguity model in which the ambiguity-adjusted weights are pooled to the "worst" prior, the

SAP model's policy lies below the MEU policy. Furthermore, as the coefficient of relative ambiguity aversion increases, the SAP policy converges to the MEU policy. Similarly, as the smooth-Choquet model is a less extreme version of CEU, the SCEU policy lies below the CEU policy.

Despite a steeper policy trajectory under the CEU model, the abatement levels in the immediate future under CEU are lower than those under MEU. This results largely due to the Choquet model's non-additive probability measure. Note that in the Choquet model, the residual probability mass, via the Choquet integral, is assigned to the worst state of the world, which is in our problem, state of the world 2015. In the immediate future, the marginal benefit of abatement is hence lower than it would be if the residual were not added. Even if the residual were not to exist, the marginal benefit of abating in 2015 would be lower in the CEU case, due to the model's assumption of the smallest hazard rate every period.

The abatement policy in CEU catches up quickly in the next decade and exceeds that of the MEU in the year 2030. After this period, abatement is notably more stringent under the CEU model. Again, we attribute this phenomenon to the non-additivity of the CEU probability measure. To minimize the residual probability mass, the agent must minimize the discrepancy between the hazard rates used to construct the survival rate (MEU hazard rates) and the hazard rate used to represent the probability of tipping. This comes down to minimizing the gap between the smallest and largest tipping probabilities under all one-step-ahead priors. *Figure 2*, shows that the discrepancies among priors are worst when atmospheric temperatures rise approximately between 1.5 and 2.5 degrees celsius, peaking at around 2 °C. Although the 1.5 °C mark (1.56 °C to be exact) is impossible to avoid exceeding, a strong near-term abatement policy under the Choquet model controls temperatures to be well below the peak discrepancy point. In the year 2095 of the "never" state of the world, the atmospheric

temperature for the MEU model is 2.146 °C, while that of the CEU model is 1.769 °C.

3.2 Likelihood of Tipping

Savage's "sure-thing principle" [Savage Leonard, 1954] is relaxed to incorporate preference for hedging for the ambiguity averse agent. However, I find that in a stochastic control setting, in which there exist two types of priors—the hazard rate (likelihood of tipping per period) and the state likelihoods (probabilities assigned to each state of the world)—, the agent's hedging strategy differs greatly based on which prior the agent deems ambiguous.

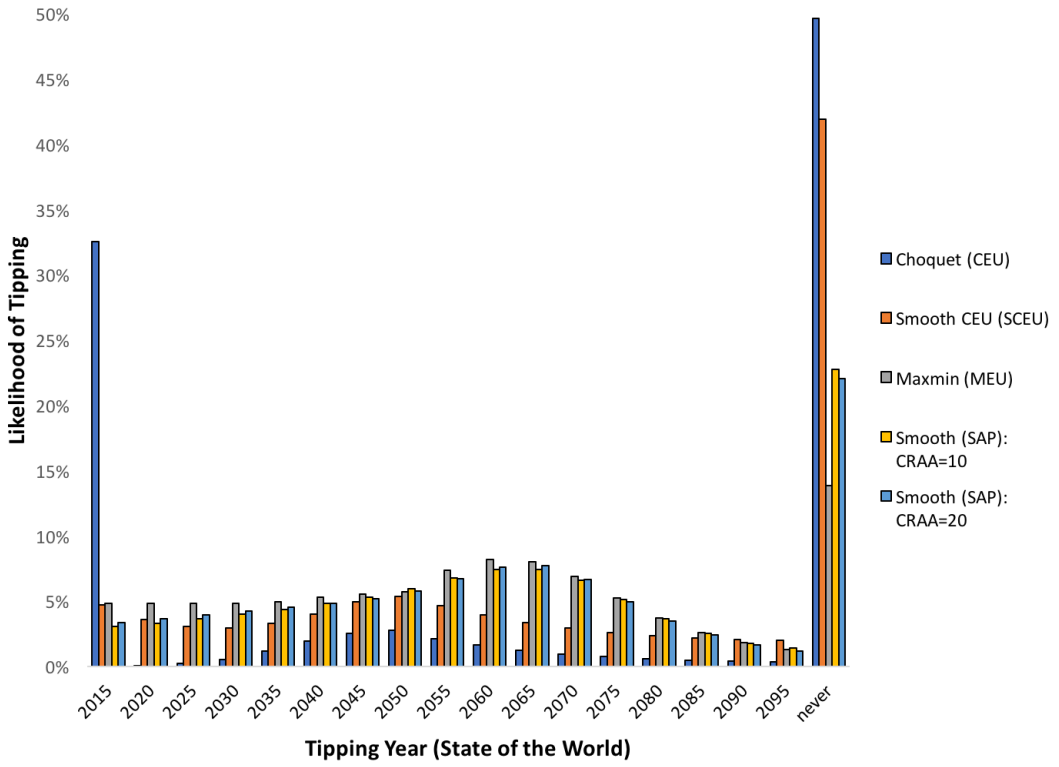

Figure 4. Likelihood of States of the World under Optimal Policies

The resulting state likelihoods under each ambiguity attitude are presented in *Figure 4*. We take away two points from this plot, again by comparing the output of CEU and MEU. Firstly, MEU places the least weight on the "never" state of the world. This should not

be surprising since at every period, the agent assumes the prior that results in the highest hazard rate. Since the weight assigned to the never tipping scenario only consists of the product of survival rates $(1 - hr)$, in essence the model is *defined* by its capacity to minimize the likelihood of the "never" state. The hedging strategy employed in MEU results in the state likelihood distribution that exhibits the least inequality. The agent hedges against the risk of tipping by assigning significant weight to all states of the world.

In contrast, I find that the agent's weight assignment under the CEU model exhibits *extreme* hedging; that is, over 75% of the weight is placed on the worst and best states of the world. Given ambiguous information on the probability of tipping at each period, and in turn, uncertainty in the likelihood of when tipping would occur, the agent hedges against the risk of climate tipping by maximizing the likelihood of the best scenario while assuming the worst hazard rates, and concurrently assuming the worst of the ambiguous information to maximize the likelihood of the worst case scenario. This results in a state probability distribution that is the most polarized among models of ambiguity.

The reason for this large discrepancy in hedging strategy between CEU and MEU is in the difference in space over which hedging is defined; that is, both weight distributions exhibit a preference for hedging, but in different dimensions. In MEU, policies, or *acts*—a mapping of states to outcomes—are defined every period over the two binary partitions—tipping hazard and tipping survival. Hedging, via a period specific abatement policy, is also defined accordingly. In CEU, acts are defined over whole consumption trajectories over the economic horizon (states of the world), characterized by when tipping takes place. In DICESC, there are 18 states of the world that include the "never" state, in which tipping never occurs. Acts in CEU are thus defined as state specific abatement trajectories, and the notion of hedging accordingly consists of assigning weights to each state of the world.

4. Risk Aversion and Pre-tipping Threshold Effects

I have thus far established how priors are assumed every period in models of multiple-prior ambiguity aversion. The hazard rate each period is defined as a function of temperature increase under a certain prior (equation (1)). It represents Bayesian learning in regards to the likelihood of climate tipping. Conditional on the event that climate tipping does not take place (the agent is still on the "never" state of the world), it also functions as a measure of how much of the underlying uncertainty is resolved per period.

The level of optimal Bayesian learning attained per unit temperature increase (henceforth referred to as "learning") is displayed in *Figure 4*. The plot shows the sensitivity of learning to the shape of the assumed prior distribution, as shown in the uniform prior and truncated normal prior learning curve. But under ambiguity aversion, optimal learning is also responsive to which priors are assumed. The latter characteristic is most evident in the CEU and MEU models—frameworks that only employ a single prior to determine the per period hazard rate—as the respective learning curves display kinks, or *points of transition* from one prior to another.

Figures 5.1 & 5.2 display the cost of a unit increase in the hazard rate, and demonstrate the transition of priors under ambiguity aversion. Under MEU, a single prior that outputs the highest hazard rate per temperature increase is assumed every period. Hence the agent first assumes the uniform prior until temperatures exceed the first transition point ($\approx 1.3\,^{\circ}\text{C}$), at which the one-sided truncated normal distribution is assumed. At the $1.56\,^{\circ}\text{C}$ mark however, a significantly more *abrupt* prior change takes place from the truncated normal prior to the normal prior. The non-smoothness of the transition is reflected in the MEU learning curve in the year 2050 (*Figure 4*).

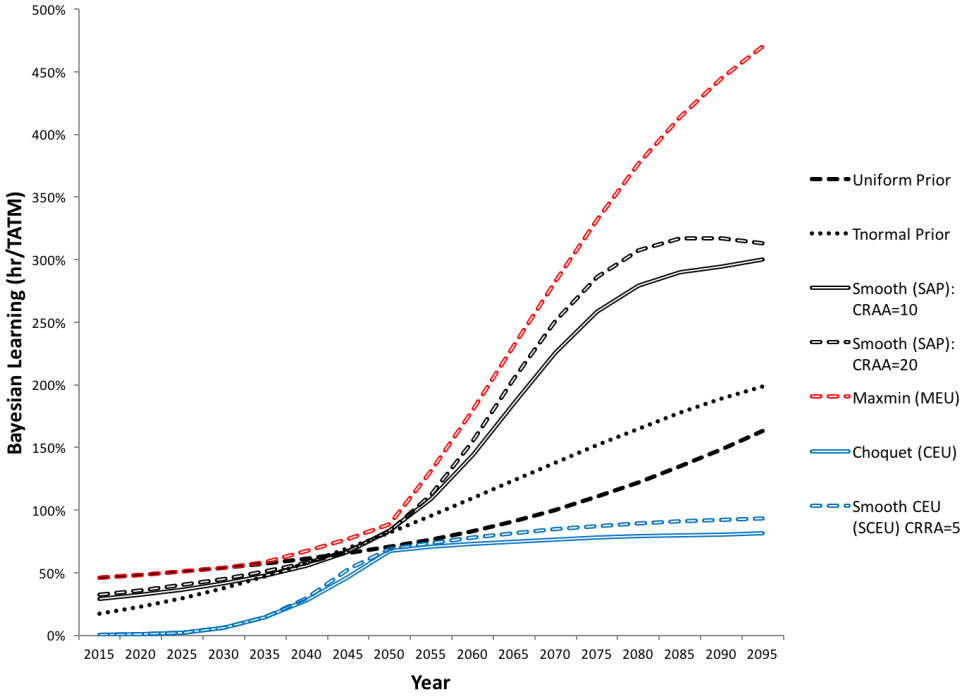

Figure 4. Optimal Bayesian Learning Across Models of Ambiguity Aversion

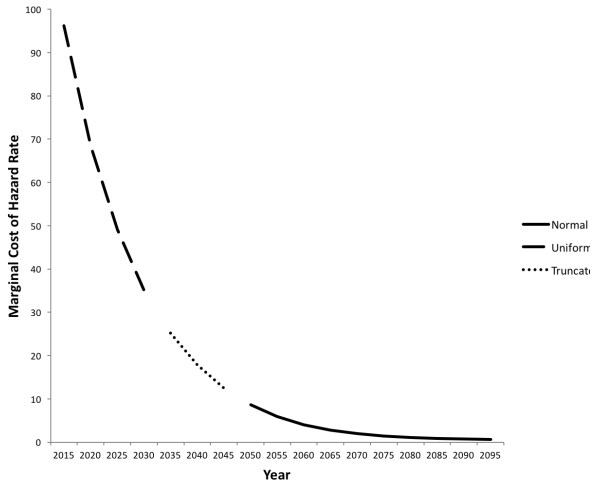

Figure 5.1 Marginal Cost of Hazard Rate: Maxmin Ambiguity Aversion (MEU)

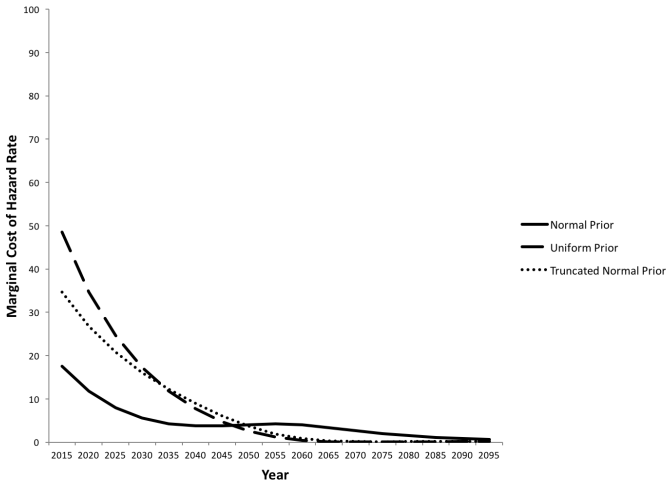

Figure 5.2 Marginal Cost of Hazard Rate: Smooth Ambiguity Aversion (SAP)

In the SAP model, the hazard rate is defined by a convex combination of all three priors, weighted by "ambiguity-adjusted weights". The marginal cost plot in *Figure 5.2* consequently shows a *weighted* marginal cost of hazard rate for each prior. The results are not too different from that of the MEU model. Notice that the priors with the highest marginal cost exactly corresponds to the priors assumed throughout time in the MEU model. I omit the graphical analysis of the marginal cost of hazard rate under the CEU model, due to the difficulty of representing the conflicting roles priors play in separately defining the hazard rate and the survival rate.

There are a couple of points to take away from *Figure 4*. First note that the Bayesian learning curve for CEU represents learning attained per unit increase in temperature when defining the hazard rate. The MEU learning curve is obtained by assuming a series of the most convenient priors—priors that obtain the lowest hazard rate each period. But also note that in constructing the survival rate, CEU assumes the worst prior every period, equivalent to the way priors are chosen in the MEU model. By focusing on the ambiguity of state likelihoods, the CEU model agent can hence apply the notion of ambiguous information to Bayesian learning. As a result, learning in the CEU model can only be characterized as being restricted to a lower (CEU) and upper bound (MEU).

Another point worth noting is that prior transitions can be abrupt, resulting in a sudden change in the pace of Bayesian learning. This is especially salient in MEU and CEU models that assume a single prior to determine the hazard and survival rate. In contrast to the per temperature learning curves of DICESC and the smooth ambiguity models, the learning curves for MEU and CEU display non-smoothness. Some prior transitions are smoother than others. For instance, the first transition in the MEU case, which takes place in year 2035, does not impose a threshold effect on learning, nor the policy. The second transition however

imposes a threshold effect on both learning and policy. Although not immediately apparent in the policy trajectory, the shape of the policy trajectory is convex-concave, with an inflection point at year 2050, the point at which the second prior transition takes place.

The policy trajectory represents a pre-tipping threshold effect with respect to the prior transition point. Before the transition, the agent abates emissions at an increasing rate, which can be attributed to three incentives. The first two incentives are discussed in Chang and Rutherford [2017]; a precautionary incentive to optimize the likelihood of near-term tipping and a damage minimizing incentive to optimize potential climate damages. The agent also has an incentive to prolong the time the economy operates under a *less hazardous* prior. With two bell-shaped priors that have most of their probability mass concentrated toward the center of the risk domain, there exists a transition point at which the prior switch leads to an abrupt increase in the rate of learning. As a result, the agent is incentivized to delay or even stop the prior transition from taking place. I denote this particular incentive, the "Pre-prior-transition (PPT) incentive".

Post-prior transition, a moderately high level of abatement is still maintained, but incentivized by the first two incentives, and the regime shift to a "worse" prior. Given the absence of the PPT incentive, and the decreasing marginal benefit of precautionary abatement over time, in some cases, the policy can weaken in the post-prior transition period. Of course, this holds true provided there does not exist another abrupt prior transition that is approachable in the near future.

The aforementioned threshold effects in the policy trajectory are more noticeable under *risk aversion*. Recall that in the CEU model, the definition of ambiguity is limited to the per-period hazard rates. In the MEU model, it is the state likelihoods that ambiguous information is defined on. The definition of risk in stochastic control, can be defined under a similar

distinction.

The agent can hedge over the outcomes corresponding to the two binary partitions every period; that is, the respective welfare under the tipping scenario and the "never" scenario. Risk aversion in this case, is applied using *conditional risk mappings*, introduced in Ruszczynski and Shapiro [2006]. If the agent chooses to hedge over entire consumption trajectories that represent states of the world, risk aversion is implemented in a *comprehensive risk measure*, also described in Ruszczynski and Shapiro [2006]. To specify aversion toward risk and ambiguity in a manner that is consistent with how the agent commits to hedging, I apply conditional risk mappings to the MEU model and comprehensive risk measures to the CEU model *.

Plots showing Bayesian learning and the optimal policy under risk and ambiguity aversion are displayed in *Figures 6.1 & 6.2*. The analysis is limited to the CEU and MEU models for simplicity. In comparison to the non-risk averse run, the MEU learning curve shows the effect of risk aversion in delaying the prior transition. The "abrupt" transition point, located at $\approx 1.56\,^{\circ}\text{C}$, is exceeded three periods later (in year 2065) than the transition time in the non-risk aversion run (year 2050) displayed in *Figure 4*. The learning curve in the CEU model nevertheless shows that the agent cannot afford to delay the point of transition, located at $\approx 1.49\,^{\circ}\text{C}$. This phenomenon is partly due to the fact that the transition point is relatively closer to the initial atmospheric temperature increase of $0.8\,^{\circ}\text{C}$, so that delaying the transition is not an option. Another factor responsible is the CEU model's assumption of the most convenient prior for the hazard rate each period. Note that the transition from the normal distribution to the uniform distribution only marginally increases the rate of learning, failing to provide the agent with enough incentive to exhibit actions of postponing the transition. In the pre-prior transition period, the optimal policy trajectory for both models exhibit close to 100% abatement levels leading up to the transition point, which demonstrates a strong

*The specifics of both risk aversion models are stated in the appendix.

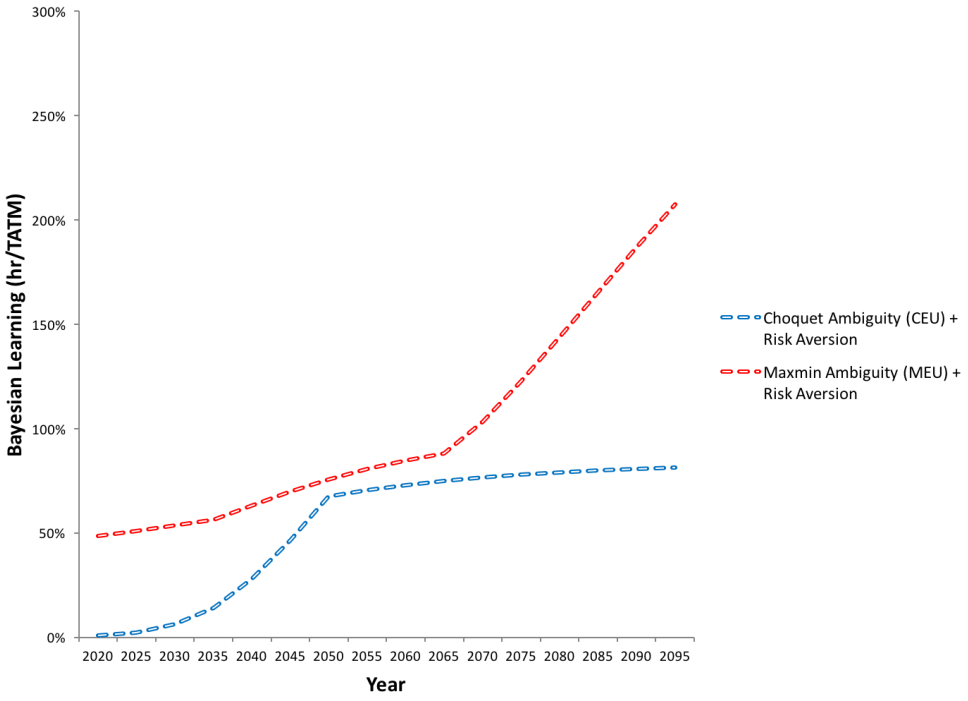

Figure 6.1 Optimal Bayesian Learning under Risk and Ambiguity Aversion

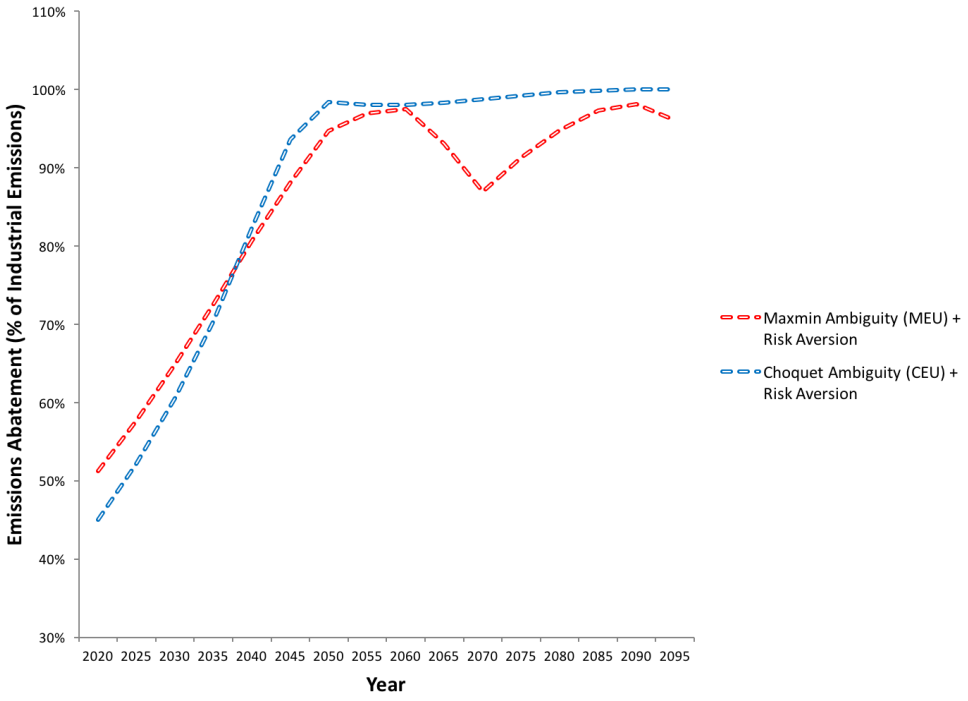

Figure 6.2 Optimal Abatement Policy under Risk and Ambiguity Aversion

incentive to delay the transition. Post transition, the absence of the PPT incentive temporarily weakens the policy by a marginal amount. The policy is reinforced shortly after, to acknowledge the increase in the rate learning due to the regime shift.

5. Conclusions

I implement four models of ambiguity aversion—including the well known Maxmin Expected Utility (MEU) model and Choquet Expected Utility (CEU) model—under the risk and uncertainty of climate tipping. Using a multiple-prior, stochastic control version of the DICE (DICESC) model introduced in Chang and Rutherford [2017], I find that aversion to ambiguity increases abatement incentives throughout the 21st century.

Results also indicate that how ambiguous information is defined matters in determining the agent's hedging strategy. If ambiguity is defined in terms of the per period hazard rate, the agent naturally hedges on a period to period basis, over the binary partition of the event tree; tipping or no tipping. In contrast, if the ambiguity is defined with respect to the probability of each state of the world taking place, the agent exhibits hedging over entire consumption trajectories. In the latter case, I find that the optimal weight assignment to states of the world exhibits extreme hedging; over 75% of the weight is placed on the worst and best states of the world.

I also identify threshold effects that determine the optimal policy before and after the transition of assumed priors takes place. In a multiple prior ambiguity aversion framework, convex combinations of the prior set can be assumed every period as a hedging strategy against the risk of tipping. If the transition of priors from one period to another can abruptly shift the world into a regime of high hazard rates, the policy exhibits an incentive to delay or even stop the prior transition from taking place. Although a moderately high level of abatement is maintained in the post-prior transition due to damage minimizing and precautionary

incentive, the absence of such "pre-prior transition incentives" can temporarily weaken the policy.

Appendix to Chapter 4: Risk Aversion

Risk Averse Stochastic Control: Risk Aversion

Recall that the utility function in DICE assumes constant elasticity of marginal utility of consumption, which is also the inverse of the *intertemporal elasticity of substitution*. In the context of the Ramsey equation, a larger value of α induces the economy to discount the future more and hence focus more on the circumstances of the present. However, the functional form

$$U[c(t)] = \frac{c(t)^{1-\alpha}}{(1-\alpha)} \tag{A1}$$

also conveniently represents constant coefficient of relative risk aversion (CRRA), as the same parameter α in $(A1)$ represents the curvature of the utility function. A greater value of α hence describes a more risk averse agent; one that puts more value on avoiding the loss of a risky outcome than the uncertain gains from gambling. This functional form is readily employed in the economics literature due to the mathematical convenience of characterizing both the concept of discounting in dynamic models and risk averse decision making with a single parameter. When discussing the risk and uncertainties of climate outcomes however, the nonseparability of the two concepts create a course for collision in our decision analysis. This is well summarized in Ackerman et al. [2013]. The functional form can only adopt two settings; one in which risk aversion is practiced at the expense of the future generation's welfare, and another that represents heightened concern for our descendants but low aversion to risky climate outcomes; none of which describe the natural environment for the optimal control problem in climate economics. Not only do the two concepts clash due to the long time-horizon of climate outcomes, but the functional form $\alpha_{risk} = 1/\alpha_{time}$ ultimately suggests parameter values that are inconsistent with observed market patterns. A stochastic

framework that can adopt separate calibration of risk aversion and time preference is thus necessary. In this appendix, I present two typs of time-risk separable risk aversion frameworks; namely, conditional risk mappings and comprehensive risk measures. I first introduce the mathematical generalization of each framework. I then propose ways to implement each framework in the DICESC model [Chang and Rutherford, 2017].

Notation

Following Ruszczynski and Shapiro [2006], Shapiro [2009] I extend the notation of the act-then learn process. I define $\mathcal{F}_t$ as the $\sigma-$algebra generated by a finite partition of set Ω_T where $\mathcal{F}_t(\xi_t)$ denotes the partition component containing $\xi_t \in \Omega_t$. The information structure is represented by the filtration $\{\mathcal{F}_t\}_t^T$ for which, $\mathcal{F}_{2015} = \{\emptyset, \Omega_T\}$ and $\mathcal{F}_{T-1}$ is the $\sigma-$subalgebra of $\mathcal{F}_T$ defined by children nodes of nodes at state $t = T-1$.

The per-period welfare function in present value terms is denoted $f_t(x_t, \xi_t)$, where f is an element of the linear space $\mathcal{Z}_t$ of $\mathcal{F}_t-$measurable functions $Z_t : \Omega_t \to \mathbb{R}$, so that for fixed x_t, $f_t(x_t, \cdot) \in \mathcal{Z}_t$. $\mathcal{Z}_t$ has dimension $|\Omega_t|$, and is $\mathcal{F}_t-$measurable. As a result, we can view $\mathcal{Z}_t$ as a subspace of $\mathcal{Z}_{t+1}$. For each t, we define a real-valued objective function:

$$\mathbb{F}_t(Z_t, \ldots, Z_T|\xi_{[1,t]}), \quad \text{for } \forall \xi_{[1,t]},\ \forall (Z_t, \ldots, Z_T) \in \mathcal{Z}_t \times \cdots \times \mathcal{Z}_T$$

which states that the objective function optimizes for future states, given past process $\xi_{[1,t]}$. The agent In DICESC, at each time period t, solves the following maximization problem:

$$\max_{x_t, \ldots, x_T} \ \mathbb{F}_t\big(f_t(x_t, \xi_t), \ldots, f_T(x_T, \xi_T)|\xi_{[1,t]}\big) \tag{P1}$$

$$\text{s.t.} \quad x_T \in \chi_\tau(x_{\tau-1}, \xi_\tau), \ \ \tau = t, \ldots, T.$$

Conditional Risk Mappings and Comprehensive Risk Measures

In DICESC, the decision maker can commit to two types of risk averse preferences that depend on different hedging strategies. The agent can hedge over binary outcomes of *(tipping, no tipping)* every period, or can also choose to hedge over states of the worlds—entire consumption trajectories what characterize the welfare under each tipping scenario. Risk aversion in the former case is implemented using *conditional risk mappings*, and in the latter, *comprehensive risk measures*. In this subsection, I present mathematical generalizations of risk aversion for each framework.

I first write out DICESC's optimization problem in terms of the notation defined previously:

$$\max_{x_{2015},\ldots,x_T} \quad \mathbb{F}\Big[f_{2015}(x_{2015}), f_{2020}(x_{2020},\xi_{2020}),\ldots,f_T(x_T,\xi_T)\Big] \tag{D1}$$
$$s.t.\ \ x_{2015}\in\chi_{2015}, x_t\in\chi_t(x_{t-1},\xi_t)\ \ t=2020,\ldots,T,$$

which can be equivalently written in nested per-period formulation $(N1)$ describing the periodic information and decision process:

$$\max_{x_{2015}\in\chi_{2015}} \quad f_{2015}(x_{2015}) + \mathbb{F}_{2020}\Bigg[\sup_{x_{2020}\in\chi_{2020}(x_{2015},\xi_{2020})} f_{2020}(x_{2020},\xi_{2020}) \tag{N1}$$
$$+\mathbb{F}_{2025}\Bigg[\ldots+\mathbb{F}_T\Big[\sup_{x_T\in\chi_T(x_{T-1},\xi_T)} f_T(x_T,\xi_T)\Big]\Bigg]\Bigg].$$

Note that formulation $(N1)$ becomes the classical risk neutral multistage stochastic programming problem when $\mathbb{F}$ is the expected value operator; i.e.

$$\mathbb{F}_t(Z_t,\ldots,Z_T|\xi_{[1,t]}) := \mathbb{E}[Z_t+\ldots+Z_T|\xi_{[1,t]}]$$

Ruszczynski and Shapiro [2006] defines risk aversion based on *conditional risk mappings*, for which a concave, monotone, *translation equivariant** mapping $\varrho_{t+1|\xi_{[1,t]}} : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$, *conditional* on history $\xi_{[1,t]}$, employs the nested formulation:

$$\max_{x_{2015}\in\chi_{2015}} f_{2015}(x_{2015}) + \varrho_{2020|\xi_{[1,2015]}}\Bigg[\sup_{x_{2020}\in\chi_{2020}(x_{2015},\xi_{2020})} f_{2020}(x_{2020},\xi_{2020}) + \varrho_{2025|\xi_{[1,2020]}}\Bigg[\ldots + \varrho_{T|\xi_{[1,T-1]}}\Big[\sup_{x_T\in\chi_T(x_{T-1},\xi_T)} f_T(x_T,\xi_T)\Big]\Bigg]\Bigg]. \tag{RA1}$$

The conditional expectation $\mathbb{E}[X_{t+1}|\mathcal{F}_t]$ is an example of a conditional mapping from χ_{t+1}, χ_t, provided that $\mathbb{E}[X_{t+1}|\mathcal{F}_t]$ is an element of space χ_t for every $X \in \chi_{t+1}$, where χ_t, χ_{t+1} are spaces of functions that are measurable by $\sigma-$subalgebras $\mathcal{F}_t, \mathcal{F}_{t+1}$, respectively. $RA1$ shows a process of risk averse hedging that maximizes the present value welfare of now and the relevant future:

$$\max_{x_{2015}\in\chi_{2015}} \underbrace{f_{2015}(x_{2015})}_{\text{Welfare Today}} + \underbrace{\varrho_{2020|\xi_{[1,2015]}}\Bigg[\sup_{x_{2020}} f_{2020} + \varrho_{2025|\xi_{[1,2020]}}\Big[\ldots + \varrho_{T|\xi_{[1,T-1]}}\big[\,\cdot\,\big]\Big]\Bigg]}_{\text{Expected Welfare in Relevant Future}} \tag{23}$$

In the context of DICESC, optimal policies optimized at initial time period 2020 and 2030 respectively, output identical abatement trajectories from 2030 onwards, conditional on $\xi^s_{2020} = \xi^{never}$.

In contrast, *comprehensive risk measures* can be written as:

$$\max_{x_{2015},\ldots,x_T} f_{2015}(x_{2015}) + \varrho_{2020}\big[f_{2020}(x_{2020},\xi_{2020})\big] + \ldots + \varrho_T\big[f_T(x_T,\xi_T)\big] \tag{RA2}$$
$$s.t.\ \ x_{2015} \in \chi_{2015}, x_t \in \chi_t(x_{t-1},\xi_t)\ \ t = 2020,\ldots,T$$

for $\rho_t : \mathcal{Z}_t \to \mathbb{R}\ \ t = 2020,\ldots,T$, which displays an optimal control process that takes into

*If $Y \in \mathcal{Z}_t$ and $X \in \mathcal{Z}_{t+1}$, then $\varrho_{t+1}(X+Y) = \varrho_{t+1}(X) + Y$

account all possibilities of $\xi_t \in \Omega_t$, regardless of the knowledge attained at time $t-1$, $\mathcal{F}_{t-1}$.

Epstein-Zin Preferences as a Conditional Risk Mapping

The Epstein-Zin utility function (henceforth denoted EZ) [Epstein and Zin, 1989] has received much attention in the stochastic DP community for climate IAMs, due to the function's separability of time and risk when applying risk aversion to the optimal policy. The recursive preference function is defined as:

$$U_t = \left[(1-\beta)c_t^{\theta} + \beta(\mu_t[U_{t+1}])^{\theta}\right]^{\frac{1}{\theta}} \tag{EZ}$$
$$\text{for } \mu_t\big[U_{t+1}\big] = \big[\mathbb{E}_t[U_{t+1}^{\alpha}]\big]^{\frac{1}{\alpha}}$$

where β is the discount rate, θ, the inverse of the intertemporal elasticity of substitution and α the coefficient of relative risk aversion. (EZ) states that current utility is a function of current welfare, plus the certainty equivalent of expected future welfare, defined by the CRRA, α. The applications of EZ preferences to the DICE model has already been implemented by a few studies in a DP framework, including the work of Jensen and Traeger [2011], Crost and Traeger [2011], Ackerman et al. [2013]. Details on implementation and the results of applying EZ to DICE or DICE-like models are well summarized in Ackerman et al. [2013]. In this section, I apply EZ preferences to DICESC in the form of a conditional risk mapping to output a time-consistent and risk averse climate policy. To maintain the conditional risk mapping form, we write the formulation as:

$$U_t = \left[(1-\beta)c_t^{\theta} + \beta(\varrho_{t+1|\xi_{[1,t]}}[U_{t+1}])^{\theta}\right]^{\frac{1}{\theta}} \tag{EZ'}$$
$$\text{for } \varrho_{t+1|\xi_{[1,t]}}[U_{t+1}] = \left[\mathbb{E}_t[U_{t+1}^{\alpha}]\right]^{\frac{1}{\alpha}}$$

where $\varrho_{t+1|\xi_{[1,t]}}$ takes the certainty equivalent present value welfare of relevant future states. Note that the mapping is a function of the period hazard rate, which is, in DICESC, controlled by the history of the abatement decision process. Depending on the realization of ξ_t^s, the conditional future utility becomes:

$$\begin{aligned} \xi_t^s = \xi^{never} \rightarrow \quad & \varrho_{t+1|\xi_{[1,t]}}[U_{t+1}] = \left[\left(1 - hr_t(\xi_{[1,t]})\right)[U_{t+1|\xi^{never}}^{\alpha}] + hr_t(\xi_{[1,t]})[U_{t+1|\xi^{t+1}}^{\alpha}]\right]^{\frac{1}{\alpha}} \\ \xi_t^s = \xi^t \rightarrow \quad & \varrho_{t+1|\xi_{[1,t]}}[U_{t+1}] = \left[U_{t+1|\xi^{t+1}}^{\alpha}\right]^{\frac{1}{\alpha}} = U_{t+1|\xi^{t+1}} \end{aligned} \tag{24}$$

Following Ackerman et al. [2013], which uses estimations of α in finance from Vissing-Jørgensen and Attanasio [2003], Bansal and Yaron [2004], I assume the value of CRRA to be $\alpha = 10$ and apply risk aversion to the distributionally robust solution (see *Figure 12*).

Nested CES Aggregation as Comprehensive Risk Measures

This approach uses separate utility aggregators over time periods and states of the world in a nested constant elasticity of substitution (CES) functional form, to represent time-risk separability in the welfare objective. This simple aggregation takes the following functional form in DICESC:

$$W = \left(\sum_s \Pi_s \left(\sum_t \beta_t C_{s,t}^{1-\theta}\right)^{\frac{(1-\gamma)}{(1-\theta)}}\right)^{\frac{1}{1-\gamma}} \tag{AD}$$

where β_t is the utility discount factor, $\alpha_{time} = \theta$, the inverse of the intertemporal elasticity of substitution used to aggregate time, and $\alpha_{risk} = \gamma$, the coefficient of relative risk aversion used to aggregate the states of the world.

Chapter 5.

Solving Stochastic Dynamic Programming Problems: a Mixed Complementarity Approach

1. Introduction

Dynamic programming (DP) is a standard tool in solving dynamic optimization problems due to the simple yet flexible recursive feature embodied in Bellman's equation [Bellman, 1957]. In the conventional method, a DP problem is decomposed into simpler subproblems characterized by a small set of state variables for which an optimal decision rule —a predefined function of the state variables —can be found at every stage. For stochastic dynamic problems in particular, DP is a powerful optimization principle, and for some stochastic problem types, DP serves as the only tractable solution method [Lontzek et al., 2012]. Methodological advances in approximation theory and numerical integration to overcome the so-called curse of dimensionality and to extend DP to problems that cannot be solved analytically, has encouraged the use of numerical DP methods, especially in economics [Judd, 1998, Rust, 1996, Maliar and Maliar, 2015, Manuelli and Sargent, 2009, Wright and Nocedal, 1999, Powell, 2011]. Complementarity programming also had a long history in economics with a formulation appearing in Scarf et al. [1967] and numerical tools introduced by rut [1995] and Ferris and Munson [2000].

Unlike the optimization algorithms in mathematical programming, there is no standard

formulation of a DP problem, nor is there an off-the-shelf solver package designated specifically for DP (as the simplex method is for linear programming problems) [Brandimarte, 2014]. DP is a principle, and there exist multiple formulations and customizations that focus on solution accuracy and computational efficiency when it comes to implementation [Tauchen, 1986, Cai and Judd, 2015, Judd et al., 2014].

Among various settings of dynamic programming problems, we focus on solving infinite horizon DP problems with continuous state and control variables. Value function iteration and policy function iteration are the two common computational approaches taken in such settings, mainly due to the algorithms' monotonic convergence properties and straightforward implementation. Despite their stability, the iterative aspect of these nonlinear programming (NLP) based algorithms make DP implementation time consuming, as processing time increases quickly in grid size for large multi-state applications, often rendering them intractable by the curse of dimensionality.

The mixed complementarity approach (MCP) introduced in this paper omits the iterative aspect of conventional NLP based DP implementations, to significantly reduce the run-time required to solve the DP problem and extend the application of dynamic programming to problems that are computationally burdensome. The MCP that corresponds to the *value function iteration* procedure is a square system of equilibrium constraints that consist of:

a. *Bellman's optimality conditions with respect to the vector of control variables, x;*

b. *optimality conditions for least squares fitting of value function parameters, α;*

for which the solution is a pair (x, α) that characterizes a Nash equilibrium with respect to the two objectives. For the latter objective, this means that in the case that the value function, $V(x\ ; \alpha)$, is estimated using polynomials, optimal parameterization of α would correspond to determining the polynomial coefficients that would obtain once the polynomial

approximation converges by the contraction mapping theorem in the iterative scheme. More importantly, computational advantages set aside, the MCP approach allows for the proper treatment of corner solutions while solving for the optimal policy [Balistreri, 1999].

Our interest in the use of complementarity methods for solving dynamic programming problems was inspired by the work of Dubé et al. [2012], Su and Judd [2012]. Their papers focus on structural estimation of discrete choice problems. Our objective is to demonstrate how their equilibrium programming approach can be extended to a continous choice problems featuring explicit complementary slackness conditions. We lay out three sample applications and argue that a mixed complementarity formulation is intuitive, robust and efficient for finding optimal policies. We stop short of showing how this approach could be applied for stuctural estimation methods, although it seems likely that Su and Judd [2012]'s ideas could be directly employed.

Following Richard Howitt's "Betty Crocker" approach to Dynamic Programming [Howitt et al., 2002a]—a pedagogic paper which makes DP more accessible to empirical economists through straightforward numerical examples—we solve various versions of the standard neoclassical growth model in GAMS to show that the one-shot MCP approach completes the value approximation procedure in a fraction of the run time required by all iterative NLP methods. We furthermore demonstrate the helpfulness of using orthogonal polynomials to improve the accuracy of the deterministic DP solution, and the use of numerical integration techniques to extend the one-shot formulation to stochastic DP problems.

The paper is organized as follows. Section 2 introduces the traditional value iteration approach in solving an infinite-horizon DP problem. Section 3 presents the corresponding DP-MCP formulation. In sections 4 and 5, we provide simple numerical examples of deterministic and stochastic optimal growth models, including a single region, 3-sector stochastic

growth model based on Global Trade Analysis Project (GTAP) data to demonstrate the computational advantages of DP-MCP. In the last section, we solve a stochastic natural resource management problem in Hydropower planning to derive optimal policies consistent with corner solutions; a key objective of the complementarity approach.

2. The Value Iteration Approach

The value function iteration procedure is the workhorse of many expository papers aimed to make numerical DP more accessible [Howitt et al., 2002a, Aruoba and Fernández-Villaverde, 2014, Manuelli and Sargent, 2009, Sargent and Stachurski, 2015]. Given an estimate of the value function, $V^n(x)$, as a function of state variable x, the value iteration approach computes an updated estimate of the value function using the Bellman equation, i.e.:

$$V^{n+1}(x) = \max_{a \in A} \left[C(x,a) + \beta V^n(x') \right]$$

where A is the action space and C, the immediate contribution function. The solution to this iterative scheme converges monotonically to the true solution of the fixed point problem, provided concavity of C and sufficient iterations of the procedure.

Traditionally based on standard nonlinear programming, the value iteration method is stable, yet due to the iterative aspect of the algorithm, is slow and time consuming. Finding all equilibria quickly becomes a daunting computational task as we increase the number state variables and grid points. In fact, computing time increases exponentially in grid size, rendering large multi-state DP applications intractable by the curse of dimensionality. As a result, complicated models that have difficulty employing recursive optimization have resorted to either simulation or math programming based stochastic frameworks [Chang and Rutherford, 2017].

The Bellman equation for the general deterministic infinite horizon DP problem with continuous state variables is stated as follows:

$$V_t(x) = \max_{a \in A(x)} \; C_t(x,a) + \beta V_{t+1}(x^{'})$$

$$s.t. \quad x^{'} = h(x,a)$$

where x is the continuous state, a, the control variable and $x^{'}$, the next-stage continuous state with transition function h.

Algorithm. ***Value Function Iteration for Infinite Horizon Problems***

1. *Set m grid points and a functional form for $V(x\ ;\alpha)$;*

   for $\forall i \leq m$, choose approximation nodes $x_i \in X$;

   fix a tolerance parameter ϵ;

   denote $V^n(x)$ to be value estimate output for iteration count n.

2. *Initialize estimate of value function $V^0(x)$.*

3. *For $n \geq 1$: obtain parameters α^{n-1} s.t. $V(x_i\ ;\alpha^{n-1}) = V^{n-1}(x_i)$*

   $\rightarrow$ *solve* $\min_{\alpha^{n-1}} \left[\sum_i \left(V^{n-1}(x_i) - V(x_i\ ;\alpha^{n-1} \right)^2 \right]$

4. *For $\forall i$, compute:*

   $V^n(x_i) = \max_{a \in A} \left[C(x_i,a) + \beta V(x_i^{'}\ ;\alpha^{n-1}) \right]$

5. *If $\|V^n - V^{n-1}\| < \epsilon(1-\beta)/2\beta$,* ***stop****;*

   else set $n = n+1$ and go to ***step*** *3.*

The algorithm seeks an approximation to the value function, such that the sum of the maximized contribution and the discounted next period value based on the approximated

function, *maximizes the total value function* [Howitt et al., 2002a]. The value iteration procedure solves for two objectives. The optimization objective in step 3 is to minimize the square deviation between estimated total value at grid points and the approximating value function. The other objective displayed in step 4, is to optimize the control variable such that the Bellman relationship holds. Convergence of the iterative approximation sequence generated by the relationship, $V^{n+1} = TV^n$, where T is the Bellman operator, is guaranteed by the contraction mapping theorem [Stokey and Lucas Jr, 1989].

3. DP-MCP: Dynamic Programming as a Mixed Complementarity Problem

We convert the value iteration process, a nonlinear optimization problem, into a nonlinear complementarity problem, a square system of equations and inequalities for which a well-established complementarity problem solver such as PATH can be applied [Ferris and Munson, 2000, Dirkse and Ferris, 1995]. The main advantage of employing a complementarity format is that it provides a simple method for incorporating inequality constraints and complementary slackness conditions. The optimization objective in our case is to find an approximation to the value function, such that the value function is maximized when Bellman's principle of optimality holds. The idea then is to solve the approximation problem simultaneously with the optimality conditions, to construct a one-shot solution to the DP problem as a complementarity problem, without resorting to value function iteration.

We note however that the complementarity formulation encompasses both primal and dual variables, doubling the number of equations. The objective of this paper is thus largely pedagogic despite the simplicity of the idea the problem reformulation is based on. The Extended Mathematical Programming (EMP) framework in GAMS functions as a useful resource for such non-standard models that require reformulation into more accessible models of established math programming classes [Ferris et al., 2009, Ferris and Sinapiromsaran, 2000].

Notice that we can write the maximization stage (step 4) of the value iteration procedure in terms of the first order conditions (FOC) of the control variable as follows:

$$\max_{a \in A} \left[C(x_i, a) + \beta V(x_i^{'} \ ; \alpha^{n-1}) \right] \longrightarrow \frac{\partial C(x_i, a)}{\partial a} + \beta \frac{\partial V(x_i^{'} \ ; \alpha^{n-1})}{\partial a} \leq 0, \ \ \forall i \in \{1, \cdots, m\}$$

Similarly, the step for function fitting by means of least-squares estimation can also be written in terms of its corresponding FOCs for $\forall n$, for which we define $\alpha = [\alpha_0, \cdots, \alpha_k]$ to be a vector of k coefficients used to parameterize the value function estimation.

$$\min_{\alpha^{n-1}} \left[\sum_i \left(V^{n-1}(x_i) - V(x_i \ ; \alpha^{n-1}) \right)^2 \right] \longrightarrow \nabla_{\alpha^{n-1}} \left[\sum_i \left(V^{n-1}(x_i) - V(x_i \ ; \alpha^{n-1}) \right)^2 \right] = 0$$

We now cast the value iteration problem as a complementarity problem by solving the maximization problem simultaneously with the least-squares function estimation problem in complementarity form. The system of inequalities and equations that make up the MCP becomes:

MCP Formulation for Deterministic Infinite Horizon DP

$$\begin{aligned}
\frac{\partial C(x_i, a)}{\partial a} + \beta \frac{\partial V(x_i^{'} \ ; \alpha)}{\partial a} \leq p_i \quad &\perp \quad a \geq 0, && \forall i \\
U_i = C(x_i, a) + \beta V(x_i^{'} \ ; \alpha) \quad &\perp \quad U_i \text{ is free}, && \forall i \\
\frac{\partial \sum_i \left(U_i - V(x_i \ ; \alpha) \right)^2}{\partial \alpha_l} = 0 \quad &\perp \quad \alpha_l \text{ is free}, && \forall l \in \{0, \cdots, k\} \\
x_i^{'} = h(x_i, a) \quad &\perp \quad p_i \text{ is free}, && \forall i
\end{aligned}$$

For the stochastic case, we can write the Bellman equation for the DP problem as follows:

$$V_t(x) = \max_{a \in A(x)} \; C_t(x, \omega, a) + \beta \mathbb{E}\{V_{t+1}(x', \omega') | x, \omega, a\}$$

$$s.t. \quad x' = h(x, \omega, a)$$

where the transition function h now includes a stochastic element $\omega \in \Omega$. The corresponding value function iteration procedure becomes:

Algorithm. ***Value Function Iteration for Stochastic Infinite Horizon Problems***

1. *Set m grid points and functional form for* $V(x, \omega \; ; \alpha)$;

   for all grid points (i, j) *choose approximation nodes* $x_i \in X$, $\omega_j \in \Omega$;

   fix a tolerance parameter ϵ;

   denote $V^n(x, \omega)$ *to be value estimate output for iteration count n.*

2. *Initialize estimate of value function* $V^0(x, \omega)$.

3. *For* $n \geq 1$: *obtain parameters* α^{n-1} *s.t.* $V(x_i, \omega_j \; ; \alpha^{n-1}) = V^{n-1}(x_i, \omega_j)$

   $\rightarrow$ *solve* $\min_{\alpha^{n-1}} \left[\sum_i \left(V^{n-1}(x_i, \omega_j) - V(x_i, \omega_j \; ; \alpha^{n-1} \right)^2 \right]$

4. *For* $\forall (i, j)$, *compute:*

   $V^n(x_i, \omega_j) = \max_{a \in A} \left[C(x_i \omega_j, a) + \beta \mathbb{E}\{V(x'_{ij}, \omega' \; ; \alpha^{n-1}) | x_j, \omega_j, a\} \right]$

5. *If* $\|V^n - V^{n-1}\| < \epsilon(1 - \beta)/2\beta$, ***stop***;

   else set $n = n + 1$ *and go to* ***step*** 3.

for which the corresponding MCP system of equations and inequalities is as follows:

MCP Formulation for Stochastic Infinite Horizon DP

$$\frac{\partial C(x_i,\omega_j,a)}{\partial a}+\beta\frac{\partial \mathbb{E}\left\{V(x_{ij}^{'},\omega^{'}\ ;\alpha)|x_i,\omega_j,a\right\}}{\partial a}\leq p_{ij}\quad\perp\quad a\geq 0,\qquad\forall(i,j)$$

$$U_{ij}=C(x_i,\omega_j,a)+\beta\mathbb{E}\left\{V(x_{ij}^{'},\omega^{'}\ ;\alpha)|x_i,\omega_j,a\right\}\quad\perp\quad U_{ij}\text{ is free},\qquad\forall(i,j)$$

$$\frac{\partial\sum\limits_{i,j}\left(U_{ij}-V(x_i,\omega_j\ ;\alpha)\right)^2}{\partial\alpha_l}=0\quad\perp\quad\alpha_l\text{ is free},\qquad\forall l\in\{0,\cdots,k\}$$

$$x_{ij}^{\prime}=h(x_i,\omega_j,a)\quad\perp\quad p_{ij}\quad\text{is free},\qquad\forall(i,j)$$

4. The Brock-Mirman Stochastic Growth Model

As a numerical exercise, we solve the Brock-Mirman Stochastic Growth Model [Brock and Mirman, 1972] compatible with Aruoba and Fernández-Villaverde [2014]. In this model, a social planner picks a consumption trajectory $\{c_t\}_{t=0}^{\infty}$ to solve:

$$\max_{c_t,k_t+1}\ \mathbb{E}\left\{\sum_{t=0}^{\infty}\beta^t\ u(c_t)\right\}$$

given utility function and resource constraint:

$$u(c_t)=\frac{c_t^{1-\theta}}{1-\theta},\quad k_{t+1}=z_tk_t^{\alpha}-c_t,$$

where β $(=0.989)$ is the discount factor, α $(=0.33)$, the capital value share and θ $(=0.95)$, the elasticity of the marginal utility of consumption. Stochastic productivity shock z_t evolves according to a 5-point Markov chain:

$$z_t\in Z=\{3.573,4.118,5.000,5.882,6.427\}$$

with transition matrix $\mathbb{Q}$ that depicts a Markov process similar to that of Tauchen [1986]'s discretized AR(1) process:

$$Q = \begin{pmatrix} 0.9727 & 0.0273 & 0 & 0 & 0 \\ 0.0041 & 0.9806 & 0.0153 & 0 & 0 \\ 0 & 0.0082 & 0.9837 & 0.0082 & 0 \\ 0 & 0 & 0.0153 & 0.9806 & 0.0041 \\ 0 & 0 & 0 & 0.0273 & 0.9727 \end{pmatrix}$$

Note that labor is fixed and the capital fully depreciates every period. We solve the model with both traditional value function iteration and DP-MCP to compare solution outputs. In both approaches, we use *Chebyshev polynomial interpolation* to estimate the value function. This particular family of orthogonal polynomials, together with an optimal choice of grid points, is widely used due to favorable convergence properties and accuracy in function interpolation. Here we briefly outline the application of Chebyshev polynomials and the first order benefits that come from it; further discussion on the application of orthogonal polynomials and Chebyshev interpolation can be found in Judd [1998].

4.1 Interpolation using Complete Chebyshev Polynomials

The complete Chebyshev polynomial with respect to the two state variables is expressed as follows:

$$V(k^{'}, z^{'} \ ; \alpha) = \sum_{0 \leq i+j \leq 2} a_{ij} T_i^c(k^{'}) T_j^c(z^{'})$$

where $T_k^c : [-1,1] \to [-1,1]$ is the k^{th} order Chebyshev basis function defined by the sinusoidal expression:

$$T_k^c(x) = \cos(k \cos^{-1}(x)), \;\; k = 0, 1, \ldots$$

Note that this closed form expression also satisfies the recursive relationship:

$$\begin{aligned} T_0^c(x) &= 1, \\ T_1^c(x) &= x, \\ T_{k+1}^c(x) &= 2xT_k^c(x) - T_{k-1}^c(x), \end{aligned} \tag{1}$$

which is almost always used in practice as it is both easier to generate and evaluate numerically than using the closed form expression. Chebyshev polynomials for given density $w(x) = (1-x^2)^{-1/2}$ are mutually orthogonal , i.e.,

$$\int_{-1}^{1} T_i^c(x) T_j^c(x) w(x) dx = 0 \;\; \forall i \neq j.$$

This orthogonality property not only makes obtaining polynomial coefficients through least-squares estimation easy for Chebyshev polynomial regressions, but also contributes to the interpolation giving the best fit over the function domain, especially when higher-order polynomials are used for function approximation.

As for determining the Chebyshev nodes, or interpolation points; using the roots of the $n+1^{th}$ order Chebyshev polynomial, as opposed to using equidistant points of the state domain, is deemed optimal in minimizing the approximation error of interpolation and in ensuring uniform convergence of the value iteration algorithm. Note that the m^{th} Chebyshev

polynomial basis function, $T_m^c(x)$, has m distinct roots on $[-1, 1]$ (i.e., $T_m^c(x) = 0$), where the roots, r_k, can be expressed as:

$$\{r_k\}_{k=1}^{m} = -\cos\left(\frac{2k-1}{2m}\right).$$

It can be shown that when the number of Chebyshev nodes is equal to the number of Chebyshev coefficients ($m = n + 1$ for interpolating polynomial of degree n) and the Chebyshev nodes are defined as the distinct m roots of the polynomial, the sup norm of approximation error is minimized [Kopecky, 2007, Mason and Handscomb, 2002].

4.2 Derivative-Friendly Recursive Definition for Chebyshev Basis Functions

The Chebyshev polynomial bases used to approximate the value function are typically defined by recursive relationship (1). In numerical DP, using this recursive relationship (as opposed to using the sinusoidal closed-form definition or any other closed-form approximation [Saad, 1980, Parlett et al., 1982]) is preferred due to the ease at which the bases are generated and numerically evaluated. However, the need to take derivatives of individual basis functions in the MCP formulation necessitates a clean, derivative-friendly expression of Chebyshev bases terms. In this subsection, we provide another way of recursively defining basis functions based on a general polynomial-form expression. This expression proves to be worthwhile especially when dealing with both first and second order derivatives of the value function when implementing shape-preserving DP-MCP in the next section.

For a k^{th} order Chebyshev polynomial, we formulate the m^{th} basis function, where $m \in \{1, \ldots, k\}$, to have $m - 1$ polynomial terms. The absolute value of the l^{th} term of the m^{th} basis

function can be expressed as:

$$2^{l-1} \cdot c_l^m x^{e_l^m} \quad \forall l \in \{1, \ldots, m-1\}$$

where x is the state variable, and c_l^m and e_l^m are respectively the polynomial coefficient and exponent used to express the l^{th} term of the m^{th} basis function.

When m is an even number, the value of the exponent, e_l^m, takes the value of the greatest odd number less than or equal to l. In contrast, when m is odd, the exponent takes the greatest even number less than or equal to l, i.e.:

m is even	m is odd
$e_1^m, e_2^m \} = 1$	$e_1^m = 0$
$e_3^m, e_4^m \} = 3$	$e_2^m, e_3^m \} = 2$
$\vdots$	$\vdots$

We can define the coefficients c_l^m recursively using the following recursive relation:

$$c_{m-1}^m = 1$$

$$c_{m-r}^m = \begin{cases} c_{m-r+1}^{m-1} & \text{if } r \in \{2, \ldots, m-1\} \text{ is even} \\ c_{m-r-1}^{m-1} + c_{m-r}^{m-2} & \text{if } r \in \{2, \ldots, m-1\} \text{ is odd} \end{cases}$$

And lastly, in writing the m^{th} basis function as a function of $m-1$ polynomial terms, the sign

of each polynomial term alternates as a function of the state variable exponent, e_l^m. Starting with the polynomial term with the highest power, which takes a positive sign, the sign of each polynomial term alternates as we decrement the exponent by two each time.

To demonstrate the construction of basis functions using the specified recursive relationship, we write out the 6^{th}-order Chebyshev polynomial basis functions in general polynomial form:

m	$l=5$	$l=4$	$l=3$	$l=2$	$l=1$
1	1				
2	x				
3	$2\cdot x^2$	-1			
4	$2^2\cdot x^3$	$-2\cdot x$	$-1\cdot x$		
5	$2^3\cdot x^4$	$-2^2\cdot x^2$	$-2\cdot 2x^2$	$+1$	
6	$2^4\cdot x^5$	$-2^3\cdot x^3$	$-2^2\cdot 3x^3$	$+2\cdot 2x$	$+1\cdot x$

Table 1. Derivative-Friendly Basis Functions

for 6^{th} order Chebyshev Polynomial

4.3 DP-MCP Formulation

We use a fourth degree Chebyshev polynomial interpolation with five interpolation points for each state variable, k_i and z_j for $i = \{1,\cdots,5\}$ and $j = \{1,\cdots,5\}$. Given lower and upper bounds $(\underline{x},\bar{x})$ on the state space, we employ linear interval conversion to map the set of interpolation points to their corresponding state space, i.e.;

$$x = \left(\frac{\bar{x}-\underline{x}}{2}\right)r + \frac{\bar{x}+\underline{x}}{2}.$$

This results in the following converted interpolation nodes for capital and productivity:

$$k_i \in \{0.122, 1.031, 2.500, 3.969, 4.878\}, \quad \text{for } k_i \in \mathbb{K} = [0, 5]$$

$$z_j \in \{3.573, 4.118, 5.000, 5.882, 6.427\}, \quad \text{for } z_j \in \mathbb{Z} = [3.5, 6.5]$$

The Bellman equation for the single sector stochastic growth model can be written as the following:

$$\begin{aligned} V_t(x) = \max_{c_t} \; & u(c) + \beta \mathbb{E}\{V_{t+1}(k^{'}, z^{'})|z\} \\ s.t. \quad & u(c) = \frac{c^{1-\theta}}{1-\theta}, \\ & k^{'} = zk^{\alpha} - c, \end{aligned}$$

where $k^{'}$ and $z^{'}$ are respectively the capital and productivity levels in the next period. For grid points (interpolation nodes) (i, j) for the two respective state variables, capital and productivity, the corresponding MCP formulation becomes:

$$\frac{\partial u(c_{i,j})}{\partial c_{i,j}} + \beta \frac{\partial \mathbb{E}\{V(k^{'}, z^{'}) \; ; \alpha)|z_j\}}{\partial c_{i,j}} \leq p_{i,j} \quad \perp \quad c_{i,j} \geq 0, \qquad \forall (i,j) \quad (2)$$

$$\frac{\partial u(c_{i,j})}{\partial k^{'}_{i,j}} + \beta \frac{\partial \mathbb{E}\{V(k^{'}, z^{'} \; ; \alpha)|z_j\}}{\partial k^{'}_{i,j}} \leq p_{i,j} \quad \perp \quad k^{'}_{i,j} \geq 0, \qquad \forall (i,j) \quad (3)$$

$$U_{i,j} = u(c_{i,j}) + \beta \mathbb{E}\{V(k^{'}, z^{'} \; ; \alpha)|z_j\} \quad \perp \quad U_{i,j} \text{ is free}, \qquad \forall (i,j) \quad (4)$$

$$\frac{\partial \sum_{i,j} \left(U_{i,j} - V(k_i, z_j \; ; \alpha)\right)^2}{\partial \alpha_l} = 0 \quad \perp \quad \alpha_l \text{ is free}, \qquad \forall l \in \{0, \cdots, 5\}$$

$$k^{'}_{i,j} = z_j k_i^{\alpha} - c_{i,j} \quad \perp \quad p_{i,j} \quad \text{is free}, \qquad \forall (i,j),$$

where α represents the vector of coefficients for the complete Chebyshev polynomial. And after discretizing the expected value operator using transition matrix $\mathbb{Q}$, we can rewrite equa-

tions $(2) - (4)$ as:

$$\frac{\partial u(c_{i,j})}{\partial c_{i,j}} + \beta \frac{\partial \sum_{j'} q(j,j') \cdot V(k',z' \ ;\alpha)}{\partial c_{i,j}} \leq p_{i,j} \quad \perp \quad c_{i,j} \geq 0, \qquad \forall (i,j)$$

$$\frac{\partial u(c_{i,j})}{\partial k'_{i,j}} + \beta \frac{\partial \sum_{j'} q(j,j') \cdot V(k',z' \ ;\alpha)}{\partial k'_{i,j}} \leq p_{i,j} \quad \perp \quad k'_{i,j} \geq 0, \qquad \forall (i,j)$$

$$U_{i,j} = u(c_{i,j}) + \beta \sum_{j'} q(j,j') \cdot V(k',z' \ ;\alpha) \quad \perp \quad U_{i,j} \text{ is free}, \qquad \forall (i,j)$$

With a tolerance level set to 1E-6, convergence of the value iteration procedure concludes in 923 iterations, close to replicating the results output by the DP-MCP formulation. Figures 1, 2 and 3 show the exponential convergence of the value iteration procedure using a five point Chebyshev interpolation, to the one-shot solution output of DP-MCP (shown using solid lines).

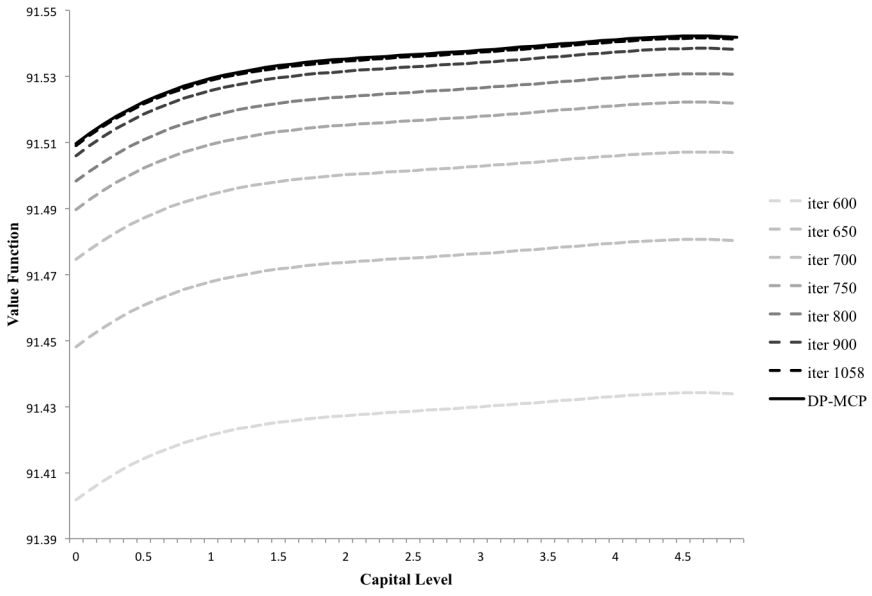

Figure 1. Value Function Approximation

iter # *plots correspond to NLP value iteration output*

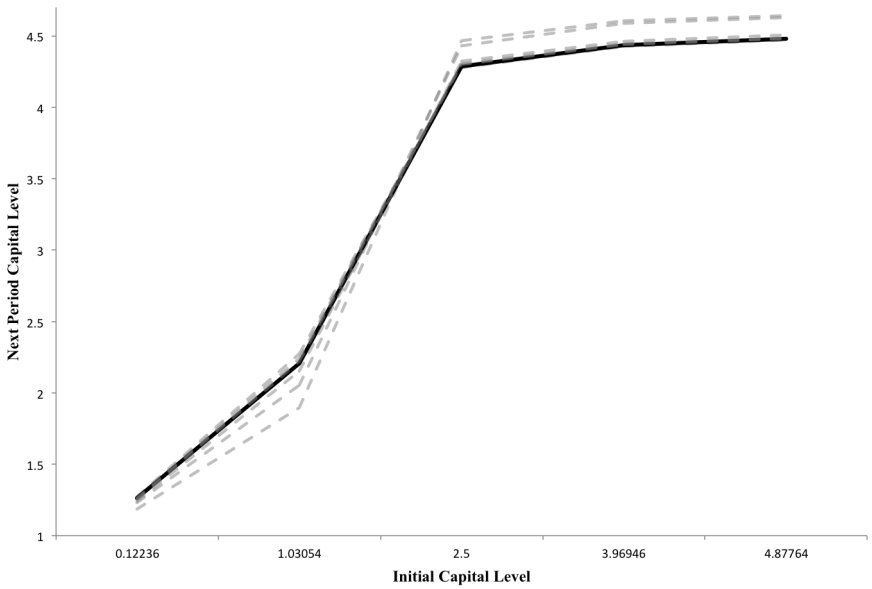

Figure 2. Subsequent Period Capital

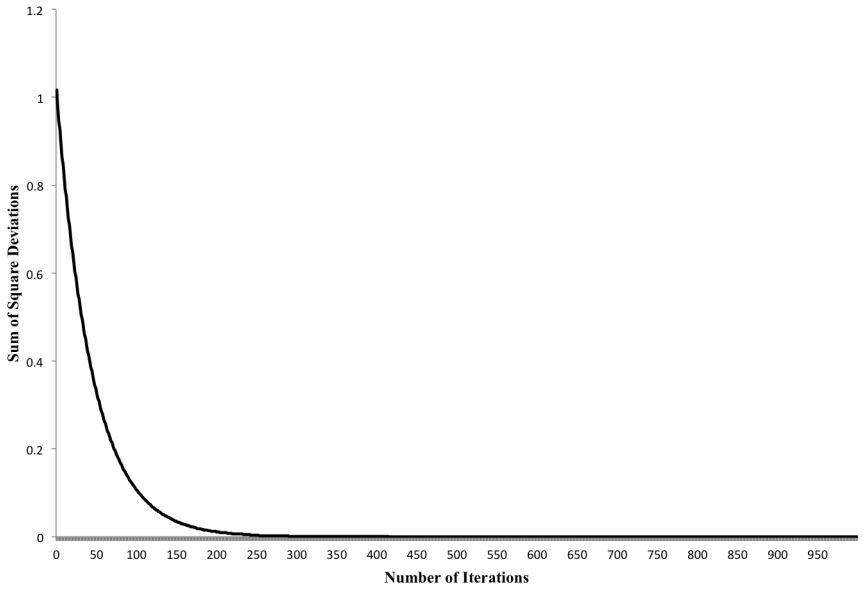

Figure 3. Convergence of Value Iteration

We report the execution times of the codes which ran on a Macbook Pro (mid 2014) with a 2.6 GHz Intel Core i5 processor with 3MB shared L3 cache and 8 GB of RAM. The average run time for the NLP value iteration procedure using the CONOPT solver in GAMS was just

above 3.5 minutes. The time for the MCP procedure to solve the same problem using the PATH solver on the other hand was on average, slightly under 0.45 seconds. We also compare our run time with the average run times of solving the same DP problem using other popular programming languages such as C++ and Fortran, as presented in Aruoba and Fernández-Villaverde [2014]. We find that solving DP-MCP with PATH on GAMS is notably faster than solving DP-NLP using any of the programs listed in the paper.

5. N-Sector Stochastic Growth Model

In this section we incorporate multiple sectors to the stochastic growth model, with perfectly correlated productivity shocks across sectors that follow a first-order autoregressive (AR1) process. A social planner chooses the consumption trajectory, $\{c_t^s\}_{t=0}^{\infty}$, investment trajectory, $\{I_t^s\}_{t=0}^{\infty}$, and labor supply trajectory, $\{L_t^s\}_{t=0}^{\infty}$, to solve the following optimization problem:

$$\max_{\{c_t^s, I_t^s, Y_t^s\}} \quad \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \;\; u(c_t^1, \cdots, c_t^n)\right]$$

such that the following contraint equations are satisfied:

- Cobb-Douglas utility function with sectoral reference consumption levels, $\bar{c}^s$, and expenditure share, η_s:

$$u(c_t^1, \cdots, c_t^n) = \prod_{s=1}^{n} \left(\frac{c_t^s}{\bar{c}^s}\right)^{\eta_s}, \quad \sum_s \eta_s = 1;$$

- the law of motion for capital accumulation with capital depreciation rate, δ $(= 0.07)$:

$$K_{t+1}^s = (1-\delta)K_t^s + I_t^s \quad \forall s;$$

- market clearing conditions for labor with labor supply fixed to the sum of all sectoral reference labor supply values:

$$\bar{L} = \sum_s L_t^s \quad \forall t;$$

- Cobb-Douglas production function with factor share of capital, α_s, productivity shock, z_t, and sectoral reference levels for output ($\bar{Y}^s$), capital ($\bar{K}^s$), and labor supply ($\bar{L}^s$). ψ_s denotes the magnitude of impact of the productivity shock to each sector:

$$Y_t^s = (1 + \psi_s(z_t - 1)) \cdot \bar{Y}_s \left(\frac{K_t^s}{\bar{K}^s}\right)^{\alpha_s} \left(\frac{L_t^s}{\bar{L}^s}\right)^{1-\alpha_s} \quad \forall s;$$

- stochastic productivity shock, z_t which follows an $AR1$ process such that the projected next period shock, $\widetilde{z}_{t+1}$ is a function of the mean productivity level, $\bar{z}$ (= 1), the persistence coefficient, ρ (= 0.9), current period productivity, z_t, and a normally distributed disturbance term, $\epsilon \sim N(0, \sigma)$ ($\sigma = 0.2$), that represents stochastic departures from the model:

$$\widetilde{z}_{t+1} = \bar{z}\left[1 + \rho\left(\frac{z_t}{\bar{z}} - 1\right) + \epsilon_{t+1}\right] \tag{5}$$

- and lastly, the equation describing the market for current output, where $a(s, ss)$ is the unit demand for intermediate good s in sector ss.

$$Y_t^s = c_t^s + I_t^s + \sum_{ss \in S} a(s, ss) Y_t^{ss} \quad \forall s.$$

The Bellman equation for the $n-$sector stochastic growth model then can be written as the following:

$$
\begin{aligned}
V_t(x) = \max_{\{c^s\}_{s=1}^n, \{I^s\}_{s=1}^n, \{L^s\}_{s=1}^n} \quad & u(c_1, \cdots, c_n) + \beta \mathbb{E}\left\{V_{t+1}(K_1^{'}, K_2^{'}, \ldots, K_n^{'}, z^{'}) | z\right\} \\
s.t. \quad & u(c_1, \cdots, c_n) = \prod_{s=1}^{n} \left(\frac{c_s}{\bar{c}_s}\right)^{\eta_s}, \quad \sum_s \eta_s = 1; \\
& K_s^{'} = (1-\delta)K_s + I^s \quad \forall s; \\
& \bar{L} = \sum_s L_s \\
& Y_s = (1 + \psi_s(z-1)) \cdot \bar{Y}_s \left(\frac{K_s}{\bar{K}_s}\right)^{\alpha_s} \left(\frac{L_s}{\bar{L}_s}\right)^{1-\alpha_s} \quad \forall s; \\
& Y_t^s = c_t^s + I_t^s + \sum_{ss \in S} a(s, ss) Y_t^{ss} \quad \forall s.
\end{aligned}
$$

Again we use a fourth degree complete Chebyshev polynomial interpolation with five interpolation points for each state variable, $k_i \in \mathbb{K} = [0, 5000]$ and $z_j \in \mathbb{Z} = [0.4, 1.6]$ for $i = \{1, \cdots, 5\}$ and $j = \{1, \cdots, 5\}$.

5.1 Discretizing the AR1 Process: Gauss-Hermite Quadrature

To evaluate the conditional expectation of the carry-over using numerical integration, we discretize the $AR1$ process using a Gauss-Hermite quadrature procedure. This step amounts to modeling the disturbance term as a random variable with Gauss-Hermite nodes as the discrete support and Gauss-Hermite weights as the corresponding probabilities. For random variable $y \sim N(\mu, \sigma)$, a Gauss-Hermite approximation of the expected value $\mathbb{E}[h(y)]$ of some function $h(\cdot)$ can be evaluated as:

$$\mathbb{E}[h(y)] = \sum_{i=1}^{m} \frac{1}{\sqrt{\pi}} \omega_i^{GH} h\big(\mu + \sigma\sqrt{2}\zeta_i\big) = \sum_{i=1}^{m} \overline{\omega}_i h(\mu + \overline{\zeta}_i), \tag{6}$$

where ζ_i and ω_i represent the original Gauss-Hermite nodes and weights. As a result, given grid points $i \in \{1, \cdots, m\}$, and disturbance term $\epsilon \sim N(0, \sigma)$, where σ is the standard deviation of the stochastic disturbance, $\overline{\zeta}_i$ then represents the *normalized* disturbance associated with Gauss-Hermite term i. The far right-hand side expression of equation (6) is ultimately used to discretize the $AR1$ process. In the context of the stochastic optimal growth problem, we discretize next period stochastic shock, $\widetilde{z}_{t+1}$, defined in equation (5) using normalized weights, $\overline{\omega}_i$, as follows:

$$\widetilde{z}_{t+1,i} = \overline{z}\left(1 + \rho\big(\frac{z_t}{\overline{z}} - 1\big) + \overline{\zeta}_i\right)$$

For the numerical examples that employ Gauss-Hermite quadratures presented in this paper, we construct a 5-point Gauss-Hermite grid with the nodes and weights specified in *Table 2.*

i	ζ_i	ω_i
1	2.0202	0.02
2	0.9586	0.3936
3	0	0.9453
4	-0.9586	0.3936
5	-2.0202	0.02

Table 2. Gauss-Hermite Approximation Data

5.2 Shape Preserving Dynamic Programming

Note that unless specified to preserve shape properties, interpolation does not assume a particular shape for the estimated value function. This can be problematic not only due to computational errors that may accumulate throughout the value iteration process, but also due to inaccuracy in the devised optimal policy. For instance, *Figure 4* displays a non-concave and even decreasing value function estimated through non-shape preserving DP formulation.

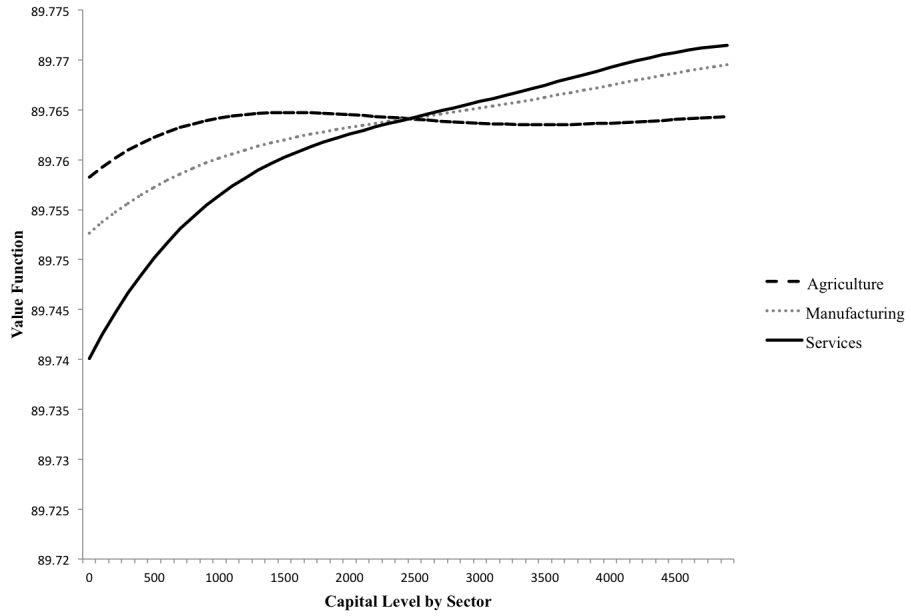

Figure 4. Approximated Value Function as a Function of Capital Levels by Sector Obtained by Non-Shape Preserving Dynamic Programming

To preserve both monotonicity and concavity of the approximated value function, we implement shaping preserving DP presented in Cai and Judd [2013]. For this we add shape constraints to DP-MCP, effectively writing the function approximation block of equilibrium

constraints as:

$$\begin{aligned}
\frac{\partial \sum_{w \in W} \left(U - V(k_w, z_w \ ; \alpha)\right)^2}{\partial \alpha_l} = 0 \quad &\perp \quad \alpha_l && \forall l \in \{0, \cdots, 5\} \\
\frac{\partial V(k_w, z_w \ ; \alpha)}{\partial k_w^s} \geq 0 \quad &\perp \quad \mu_w^{k^s} && \forall s, \forall w \in W \\
\frac{\partial^2 V(k_w, z_w \ ; \alpha)}{\partial z_w^2} \leq 0 \quad &\perp \quad \mu_w^z && \forall w \in W,
\end{aligned}$$

where W represents all states of the world characterized by a combination of discretized levels of state variables, which in our case corresponds to capital levels for each sector, K_s and productivity level z. We omit the full MCP formulation for conciseness.

5.3 Numerical Example

As a numerical example, we solve a three sector stochastic optimal growth model. The description of sectors and model parameters is summarized in *Tables 3 & 4*. The estimated coefficients of the 70 polynomial terms obtained by running DP-MCP is presented in *Table 5.*

Parameter	Agriculture	Manufacturing	Services
η_s	0.006	0.154	0.810
α_s	0.532	0.326	0.300
ψ_s	0.249	0.318	0.184
$\bar{z}_s$	4.007	5.116	2.966
$\bar{c}^s$	83.829	2061.481	10827.835
$\bar{K}^s$	109.634	836.317	3633.529
$\bar{L}^s$	96.445	1729.073	8478.235
$\bar{Y}^s$	413.749	6981.547	19502.889

Table 3. Sector Specific Reference Values

$a(s, ss)$	Agriculture	Manufacturing	Services
AGR	46.776	211.546	33.594
MFR	74.885	2651.269	1668.552
SER	100.181	1333.579	5181.908

Table 4. Intermediate Demand (s) in sectors (ss)

Coefficient	Value	Coefficient	Value	Coefficient	Value	Coefficient	Value
cp1	89.735230	cp19	-0.000440	cp37	-0.000520	cp55	-0.000120
cp2	0.003430	cp20	-0.001190	cp38	0.000490	cp56	-0.009510
cp3	-0.001780	cp21	-0.000060	cp39	-0.000130	cp57	0.000070
cp4	0.001060	cp22	0.000030	cp40	-0.001310	cp58	-0.000030
cp5	-0.000260	cp23	0.000630	cp41	0.000110	cp59	0.000150
cp6	0.007840	cp24	-0.000010	cp42	-0.000080	cp60	-0.000040
cp7	-0.000400	cp25	-0.000350	cp43	0.000340	cp61	-0.000060
cp8	0.000360	cp26	-0.004450	cp44	-0.000100	cp62	0.000130
cp9	-0.000080	cp27	0.000050	cp45	-0.000150	cp63	-0.000040
cp10	-0.001880	cp28	-0.000070	cp46	-0.001450	cp64	-0.000090
cp11	0.000120	cp29	-0.000020	cp47	0.000220	cp65	-0.000030
cp12	-0.000150	cp30	-0.000080	cp48	-0.000200	cp66	0.000390
cp13	0.000960	cp31	-0.000140	cp49	0.000180	cp67	-0.000010
cp14	-0.000060	cp32	0.001980	cp50	-0.000010	cp68	-0.000050
cp15	-0.000350	cp33	-0.000030	cp51	-0.000250	cp69	-0.000030
cp16	0.015400	cp34	-0.000120	cp52	0.000240	cp70	-0.000220
cp17	-0.001370	cp35	-0.000810	cp53	-0.000090		
cp18	0.001030	cp36	0.187090	cp54	-0.000130		

Table 5. Coefficients of 4-Dimensional Complete Chebyshev Polynomial

We display the optimal policy output by DP-MCP, by solving t sequential optimization

problems given the estimated value function (*Figure 5.*) and the simulated sequence of the productivity event process. Given a 200 year time horizon $(t = 200)$, the simulated productivity shock that follows an AR1 process is displayed in *Figure 6.* The optimal trajectory for consumption, output and capital accumulation is displayed in *Figures 7-9*.

The average time for DP-MCP to run with PATH in this case was on average, 42 minutes. The process included 150 NLP value function iterations using the CONOPT solver to provide a good starting point for DP-MCP to conclude the convergence process. We note that the iterative NLP procedure took up most of the run-time, while the MCP took less than 5 minutes to solve. Given the rate of convergence displayed by NLP iterations, we estimate at least 800 more iterations are necessary for the sum of squared errors to reach a tolerance level of 1E-6. We also note that the implementation of shape-preserving DP proved to be highly time intensive. Although shape properties were preserved only for a single sector (agriculture), the run-time increased by four fold compared to solving the growth problem using regular DP, which took less than 10 minutes (50 NLP iterations) to solve.

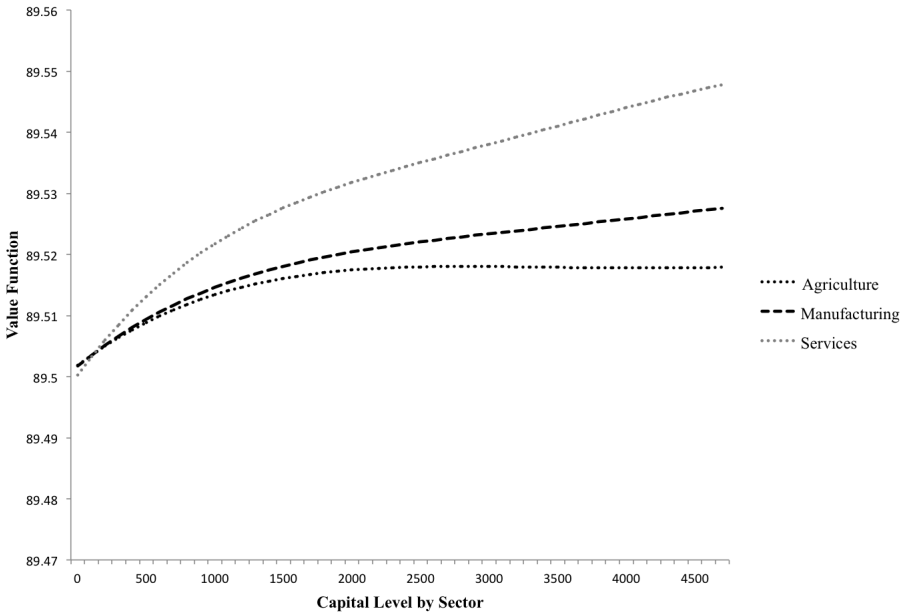

Figure 5. Approximated Value Function as a Function of Capital Levels by Sector: Shape Preserving Dynamic Programming

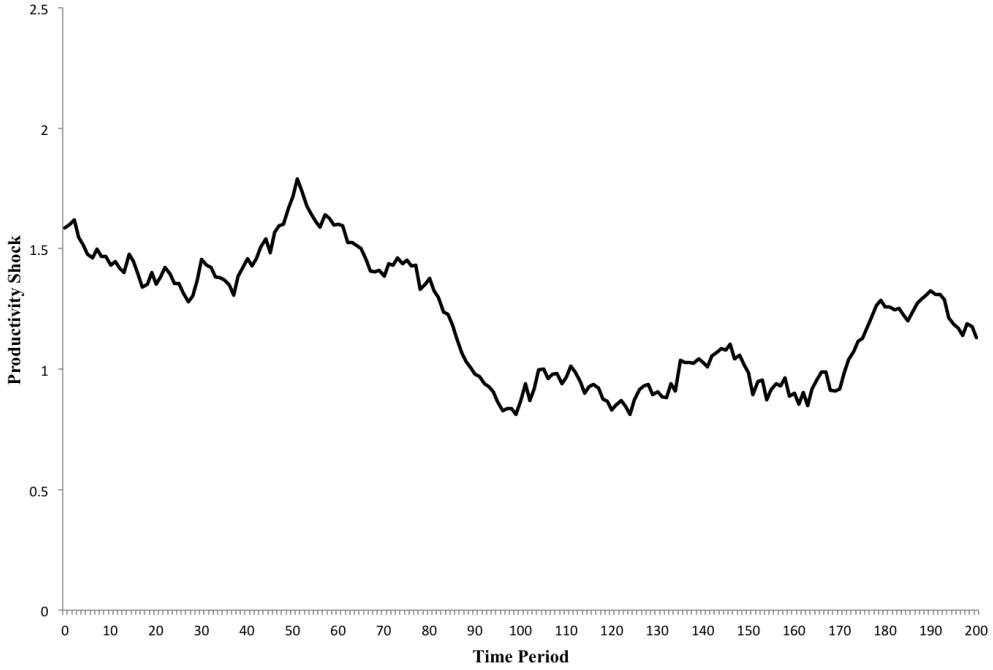

Figure 6. Simulation Results for Stochastic Productivity Levels

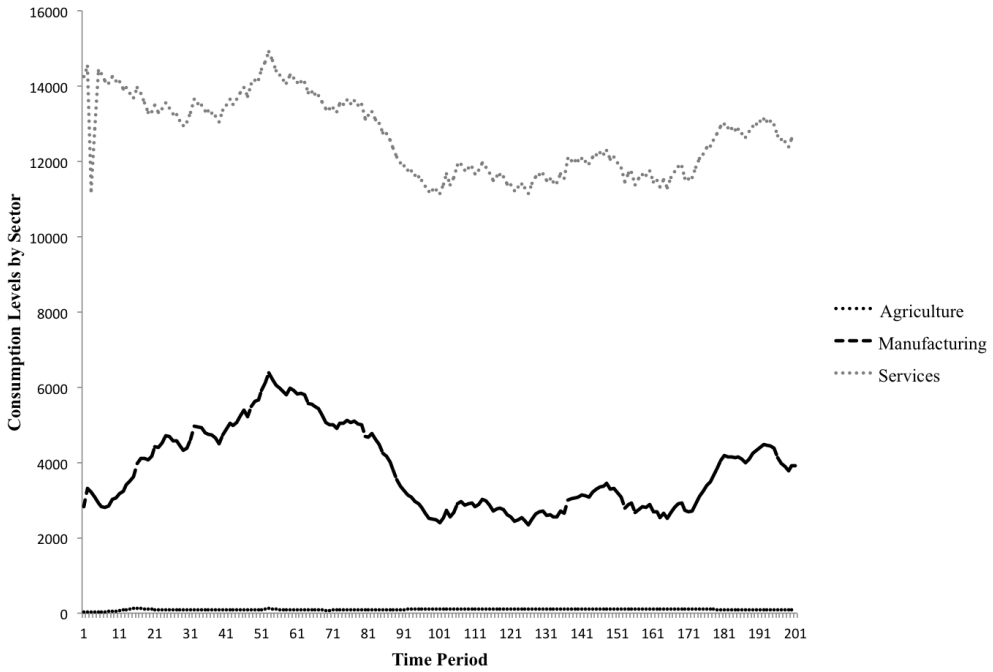

Figure 7. Simulation Results for Consumption Levels by Sector

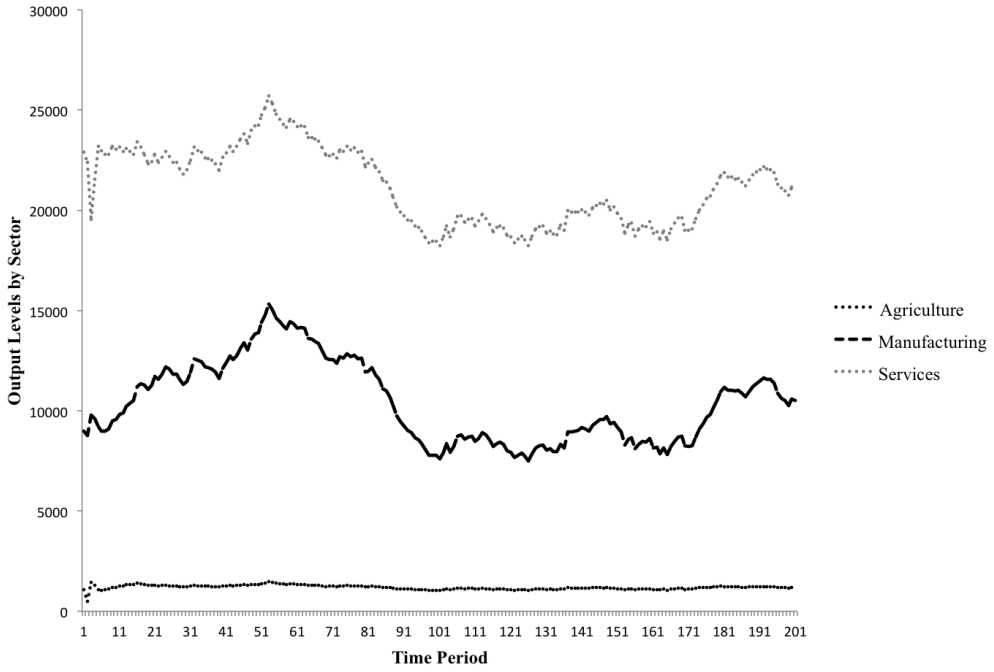

Figure 8. Simulation Results for Output Levels by Sector

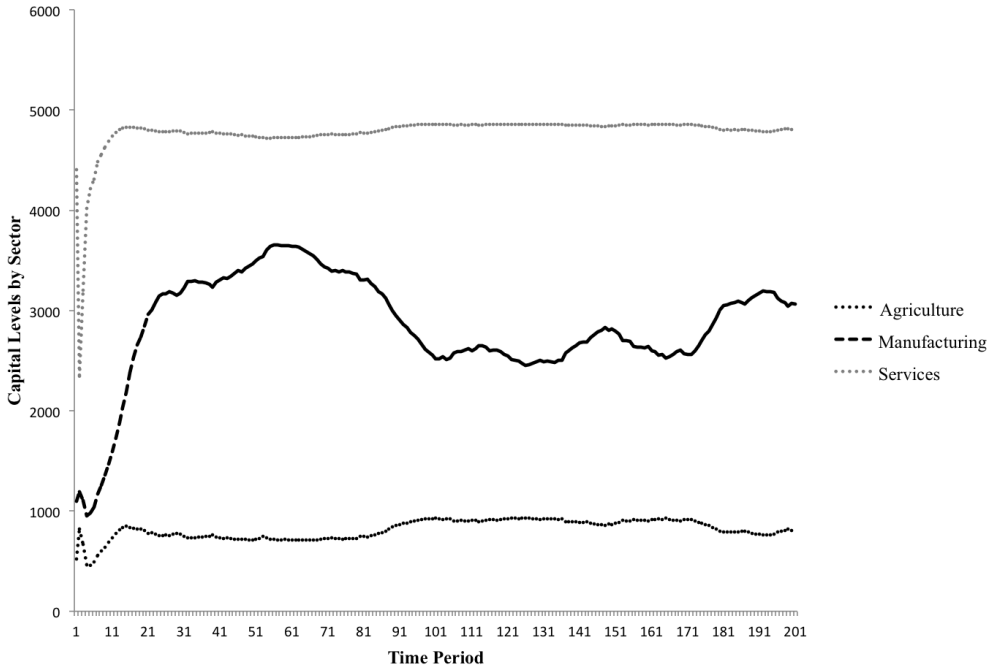

Figure 9. Simulation Results for Capital Accumulation Levels by Sector

6. Stochastic Hydropower Planning

In this section we solve a mixed complementarity problem and illustrate DP-MCP's accommodation of corner solutions. We also present quadrature methods that can be utilized in the MCP formulation to solve stochastic dynamic programming models. The hydropower planning problem is an annual model of a single aggregated reservoir with a monthly release schedule for hydropower generation. The present model is similar in dynamics to the model of water management on the North Platte River in Nebraska presented in Howitt et al. [2002a,b]. We begin with a brief overview of the model before describing algorithmic issues.

Month	Inflow
Jan	1.2
Feb	1.0
Mar	1.1
Apr	1.2
May	40.2
Jun	99.5
Jul	146.3
Aug	138.2
Sep	70.7
Oct	11.7
Nov	2.3
Dec	1.5

Table 7. Mean Monthly Inflow to Reservoir (million cubic meters)

Two state variables characterize water management at time t, namely L_t, the stock of water in the reservoir and D_t, the inflow of water to the reservoir as a function of precipitation levels. Inflow D_t consists of stochastic deviations from average monthly inflow levels (displayed in

Table 7.) characterizing the state variable as following a first order autoregressive process. z_t and r_t denote the water retained and released every month to generate electricity; if the capacity of the reservoir is exceeded, the excess water runs down the spillway and bypasses the power plant. In other words, spills (denoted s_t) balance the system when the reservoir overflows, but have no economic value in the model. The maximum capacity of the reservoir dam is 250 million cubic meters (MCM). There exists no lower bound on the water level that must be maintained in the reservoir.

Lastly, monthly electricity generation must meet a fixed monthly demand of 140 megawatt hours (MWH). In case electricity generated by hydropower does not meet demand, non-hydro electricity generation is employed, incurring marginal cost equal to the market price of electricity. We begin by writing down the primal model.

Primal Equations

Water at the start of the month is either retained or released to generate electricity:

$$L_t \geq z_t + r_t$$

Total generation of electricity through hydro (r_t) and non-hydro (x_t) sources must equal the demand for electricity. Demand (denoted g_t) is fixed in each month:

$$r_t + x_t = g_t$$

The projected level of water at the start of the subsequent month ($\widetilde{L}_{t+1}$) depends on how much water is currently stored (z_t), how much inflow is projected to occur ($\widetilde{D}_{t+1}$) and how much

water will be spilled ($\widetilde{s}_{t+1}$). Projected variable levels are represented using a tilde:

$$\widetilde{L}_{t+1} = z_t + \widetilde{D}_{t+1} - \widetilde{s}_{t+1}$$

The price of water in the subsequent month ($\widetilde{p}_{t+1}$) is imputed on the basis of the imputed water price (value of water as a function of the state):

$$\widetilde{p}_{t+1} = V\big(\widetilde{L}_{t+1}, \widetilde{D}_{t+1} \ ; \alpha\big) \tag{7}$$

Lastly, inflow D_t follows an $AR1$ process such that the projected inflow $\widetilde{D}_{t+1}$ is a function of the mean monthly inflow $\overline{D}_t$, the coefficient for rainfall persistence, ρ, the realized inflow levels, D_t, and a normally distributed disturbance term $\epsilon \sim N(0, \sigma)$ that represents stochastic departures from the model.

$$\widetilde{D}_{t+1} = \overline{D}_{t+1}\left[1 + \rho\left(\frac{D_t}{\overline{D_t}} - 1\right) + \epsilon_{t+1}\right] \tag{8}$$

The objective is to minimize non-hydro electricity generation while meeting the fixed demand for electricity. This is presented in the following objective function:

$$-\overline{c}\overline{x}\left(\frac{x_t}{\overline{x}}\right)^{\eta} + \beta\mathbb{E}\left\{\widetilde{p}_{t+1}\widetilde{L}_{t+1}|D_t\right\},$$

where $\overline{c}$ is the reference cost of non-hydro generation, $\overline{x}$ is the reference supply of non-hydro generation and η, the elasticity of non-hydro supply. Using the value function expression in (7) to estimate the projected shadow price of water, $\widetilde{p}_{t+1}$, we can rewrite the objective as follows:

$$\underbrace{-\overline{cx}\left(\frac{x_t}{\overline{x}}\right)^{\eta}}_{contribution} + \underbrace{\beta\,\mathbb{E}\left\{\left(V\left(\widetilde{L}_{t+1},\widetilde{D}_{t+1}\ ;\alpha\right)\right)\widetilde{L}_{t+1}|D_t\right\}}_{carry-over}.$$

To discretize the $AR1$ process, we again use a 5-point Gauss-Hermite quadrature method. Given normalized Gauss-Hermite disturbance $\overline{\zeta}_i$, associated with normalized weight, $\overline{\omega}_i$, we can rewrite equation (8) as follows:

$$\widetilde{D}_{t+1,i} = \overline{D}_{t+1}\left(1+\rho\left(\frac{D_t}{\overline{D}_t}-1\right)+\overline{\zeta}_i\right) \quad \forall i \in \{1,\ldots,5\}.$$

We write down the corresponding Bellman equation to the primal problem. For conciseness of notation, we drop all time subscripts. Note again that the tilde sign is used to denote *imputed* next period variables.

$$\begin{aligned}
\max_{r,z,x}\ & -\overline{cx}\left(\frac{x}{\overline{x}}\right)^{\eta} + \beta\sum_i \overline{\omega}_i\left[V\left(\widetilde{L}_i,\widetilde{D}_i\ ;\alpha\right)\right]\widetilde{L}_i, \quad s.t. \\
& L \geq z + r \\
& r + x = g \\
& \widetilde{L}_i = z + \widetilde{D}_i - \widetilde{s}_i \quad \forall i \\
& \widetilde{D}_i = \overline{D}_i\left[1+\rho\left(\frac{D}{\overline{D}}-1\right)+\overline{\zeta}_i\right] \quad \forall i
\end{aligned}$$

Dual Equations

We now state the dual variables and equations used for the MCP formulation.

FOC wrt x_t The zero profit condition on non-hydro electricity generation, which states that the marginal cost of non-hydro generation equals the market price of electricity (p_t^e):

$$\bar{c}\eta \frac{x_t{}^{(\eta-1)}}{\overline{x}} \geq p_t^e$$

FOC wrt r_t The value of electricity obtained through hydro-generation is equal to the shadow value of water:

$$p_t \geq pe_t$$

FOC wrt z_t Water left in the reservoir equals the expected value of water in the subsequent month:

$$p_t \geq \beta \sum_i \overline{\omega_i} \left[\frac{\partial V\left(\widetilde{L}_{t+1,i}, \widetilde{D}_{t+1,i} \ ; \alpha\right) \widetilde{L}_{t+1,i}}{\partial z_t} \right]$$

FOC wrt $\widetilde{s}_{t+1,i}$ Spilling water amounts to free disposal. This assures that the shadow value of water in subsequent months is nonnegative:

$$\frac{\partial V\left(\widetilde{L}_{t+1,i}, \widetilde{D}_{t+1,i} \ ; \alpha\right) \widetilde{L}_{t+1,i}}{\partial \widetilde{s}_{t+1,i}} \geq 0$$

MCP Formulation

We choose five Chebyshev interpolation nodes for the reservoir water level based on the

minimum and maximum water levels permitted for the operation of the dam; i.e.,

$$L^j \in \mathbb{L} = \{6.118, 51.527, 125, 198.473, 243.882\}, \;\; j = \{, 1, \cdots, 5\}$$

Similarly, the nodes for water inflow, D^k, are determined by the average monthly inflow $\overline{D}$ and the normalized standard deviation of inflow σ such that: $D^k \in [\overline{D} - 3\sigma, \overline{D} + 3\sigma], \;\; k \in \{1, \cdots, 5\}$.

We begin by specifying the *value maximization* block of equations that consist of the first order optimality conditions for control variables within the MCP. All equations and variables are now indexed with grid point indices. For conciseness of notation, we omit the time subscripts.

$$\begin{aligned}
L^j \geq z_{j,k} + r_{j,k} \quad &\perp \quad p \geq 0, && \forall (j,k) \\
r_{j,k} + x_{j,k} = g_{j,k} \quad &\perp \quad p^e \geq 0, && \forall (j,k) \\
\widetilde{L}_{j,k,i} = z_t + \widetilde{D}_{j,k,i} - \widetilde{s}_{j,k,i} \quad &\perp \quad \widetilde{L}_{j,k,i} \text{ is free}, && \forall (j,k,i) \\
\widetilde{p}_{j,k,i} = V\big(\widetilde{L}_{j,k,i}, \widetilde{D}_{j,k,i} \;; \alpha\big) \quad &\perp \quad \widetilde{p}_{j,k,i} \text{ is free}, && \forall (j,k,i) \\
\bar{c}\eta \frac{x_{j,k}}{\overline{x}}^{(\eta-1)} \geq p^e_{j,k} \quad &\perp \quad x_{j,k} \geq 0, && \forall (j,k) \\
p_{j,k} \geq p^e_{j,k} \quad &\perp \quad r_{j,k} \geq 0, && \forall (j,k) \\
p_{j,k} \geq \beta \sum_i \overline{\omega_i} \left[\frac{\partial V\big(\widetilde{L}_{j,k,i}, \widetilde{D}_{j,k,i} \;; \alpha\big) \widetilde{L}_{j,k,i}}{\partial z_{j,k}} \right] \quad &\perp \quad z_{j,k} \geq 0, && \forall (j,k) \\
\frac{\partial V\big(\widetilde{L}_{j,k,i}, \widetilde{D}_{j,k,i} \;; \alpha\big) \widetilde{L}_{j,k,i}}{\partial \widetilde{s}_{j,k,i}} \geq 0 \quad &\perp \quad \widetilde{s}_{j,k,i} \geq 0, && \forall (j,k,i)
\end{aligned} \tag{9}$$

The value function is estimated using a 4^{th} order complete Chebyshev polynomial with

respect to the two state variables, L_k and D_k. Optimality conditions for the least-squares value function fitting are specified as follows:

$$\frac{\partial \sum_{j,k} \left(p_{jk} - V(L_j, D_k \ ; \alpha)\right)^2}{\partial \alpha_l} = 0 \quad \perp \quad \alpha_l \text{ is free}, \quad \forall l \in \{1, \cdots, 15\} \tag{10}$$

Implementing DP-MCP amounts to solving the system of equations and inequalities in (9) and (10). The approximated shadow price of water and optimal release schedule is displayed in *Figures* $10 - 13$. Each plot corresponds to an approximation node for reservoir water levels. The high and low scenarios of precipitation correspond to inflow levels that are three standard deviations higher and lower than the monthly average respectively. The resulting shadow price of water is high when both the stock of water and inflow are low, and is equal to zero during the rainy months especially when the reservoir is sufficiently filled. The optimal release schedule displays the opposite dynamics as anticipated.

Lastly, given the estimated value function, we run a five year simulation of stochastic inflows for which we solve for the optimal water release (hydropower generation) trajectory. The optimal release schedule is displayed in *Figure 14.*

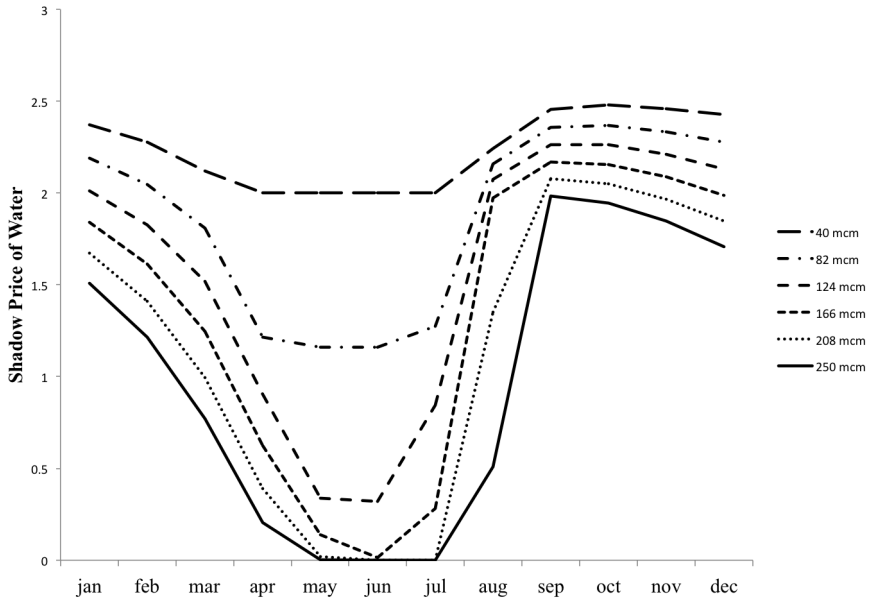

Figure 10. Shadow Price of Water Approximation for High Precipitation Scenario

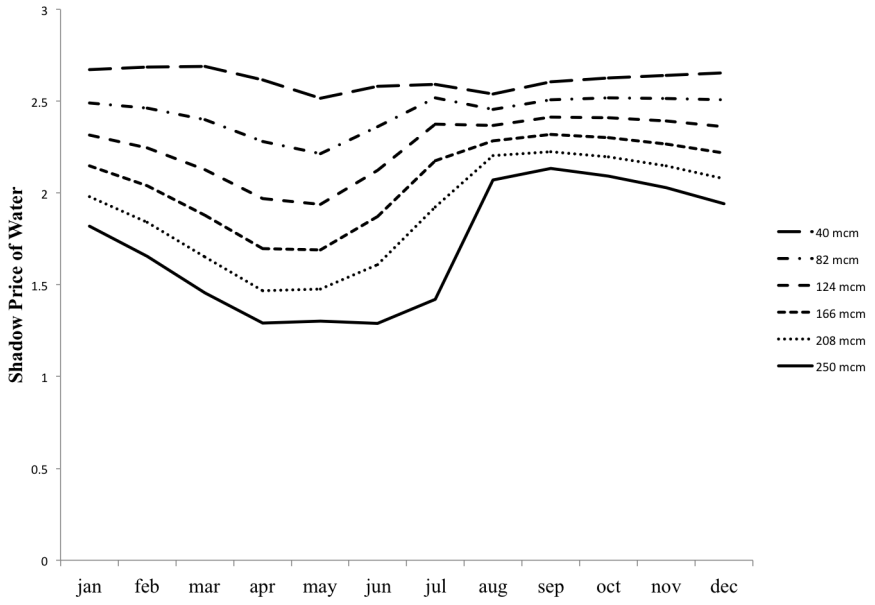

Figure 11. Shadow Price of Water Approximation for Low Precipitation Scenario

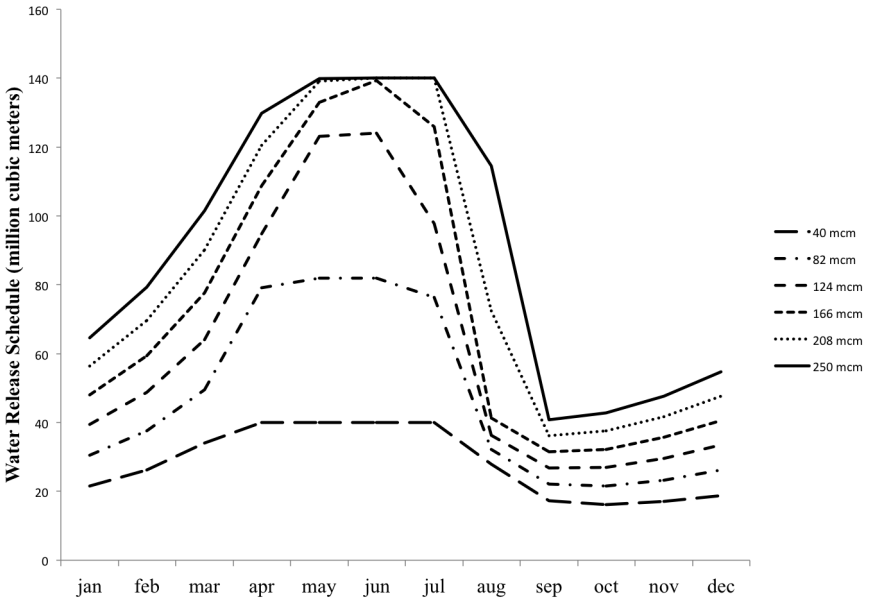

Figure 12. Release Schedule for High Precipitation Scenario

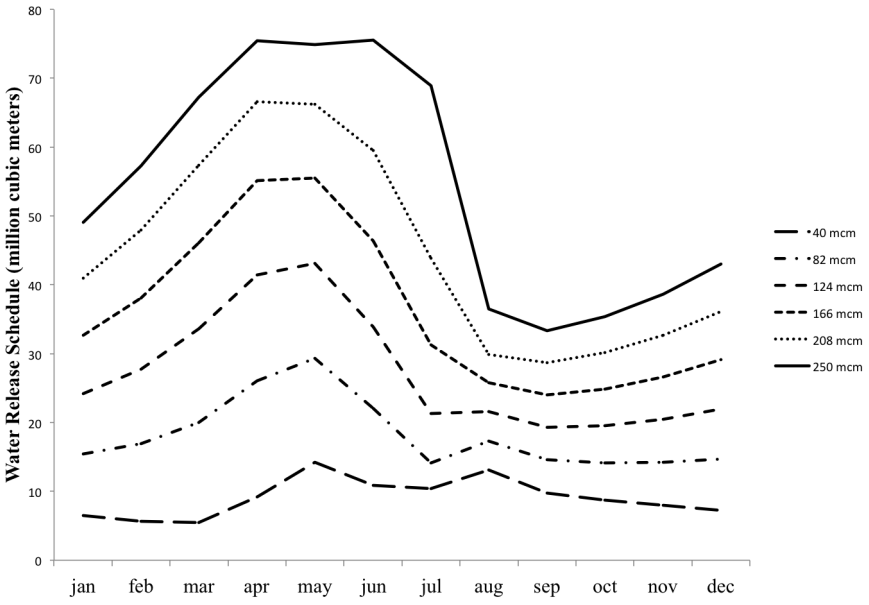

Figure 13. Release Schedule for Low Precipitation Scenario

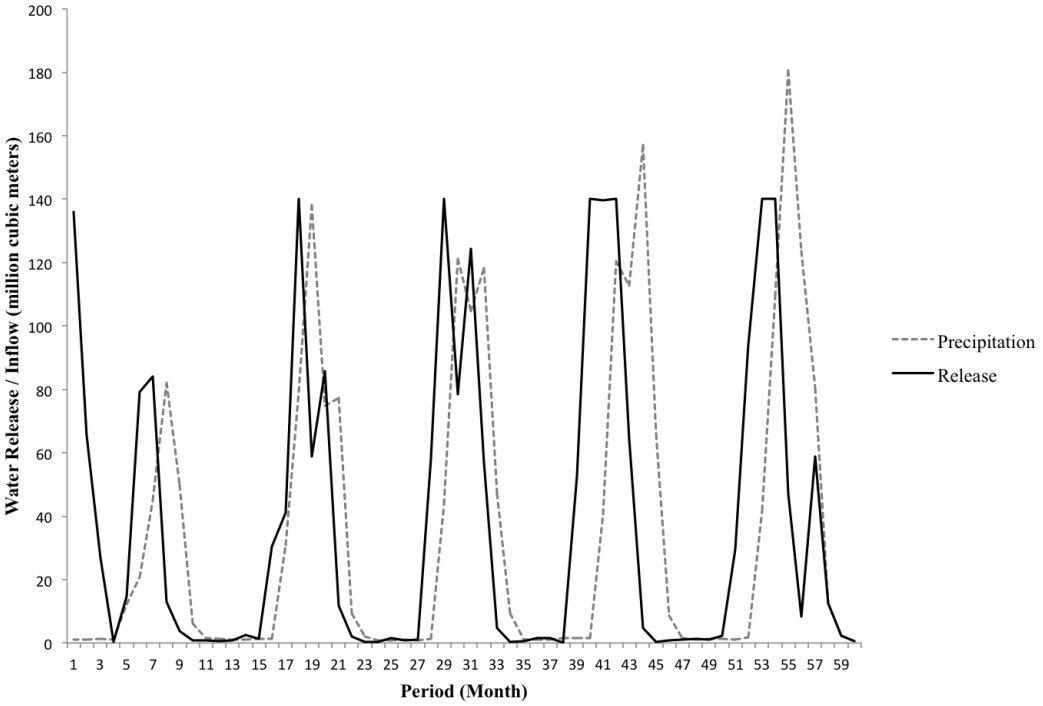

Figure 14. Simulation Result for Release Schedule

7. Conclusions

In this paper, we introduce a mixed complementarity approach (DP-MCP) for solving infinite horizon, continuous-state dynamic programming problems using off-the-shelf MCP solvers. We benchmark the value function iteration procedure which is traditionally based on iterative NLP methods, and solve a system of equilibrium constraints that characterize a Nash equilibrium in the two subproblems of the value iteration algorithm; namely, *a.* finding state and control variables in accordance with Bellman's principle of optimality, conditional on the estimate of the value function; and *b.* optimizing the parameters of the value function estimate given the vector of optimized state variables. Using numerical examples, we demonstrate that the oneshot DP-MCP approach significantly reduces the run-time required to solve the

DP problem by means of eliminating the iterative aspect of NLP based DP implementations. We further stress the computational advantages by extending the application of dynamic programming to stochastic dynamic problems that are computationally burdensome for conventional NLP based approaches to process. More importantly, the MCP approach allows for the proper treatment of corner solutions via complementary slackness conditions while solving for the optimal policy, making the complementarity formulation of DP both efficient and robust.

Appendix to Chapter 5: GAMS Code

1. Brock-Mirman Stochastic Growth Model

```
$title Stochastic Brock-Mirman Growth Model Using Chebyshev Polynomial Approximation:

*     This program includes both DP-NLP and DP-MCP formulations

Sets  s     State variables                          /cap,phi/,
      ik    Nodes for capital at which value function is evaluated      /1*5/,
      ip    Nodes for productivity                          /1*5/,
      ic    Dimension of Chebyshev polynomial               /1*5/,
      iter  Dynamic programming iterations                        /1 * 2000/;

alias (ik,jk)
alias (ip,jp)
alias (ic,jc)
alias (s,ss)

Parameters
      eta   Elasticity of the marginal utility of consumption /0.95/,
      beta  Utility discount factor /0.9888/,
      alpha Capital value share     /0.333/,
      pi    /3.141593/;


*     Parameters to define CS Polynomial terms
*     Defined for both capital (K) and productivity (p)

Parameters
      arg_k, arg_p      Argument of cosine weighting function,
      x_k, x_p          Node value for the state variable on the unit interval,
      lo_k, lo_p        Lowerbound on stock variable,
      up_k, up_p        Upperbound on stock variable,
      csbar(ic)         Chebyshev polynomial terms,
      cap(ik)           Stock level value at node for grid point calculation,
      phi(ip)           Stock level value at node for grid point calculation;

Parameters
      sigma             Normalized standard deviation of phi /0.1/,
      p_mean            Mean value of productivity phi      /5/,
      p_std             Standard deviation of phi;


*     Set lower and upper bound on state variables

p_std = sigma * p_mean;
lo_p = p_mean - 3*p_std;
up_p = p_mean + 3*p_std;

lo_k = 0;
up_k = 5;

*     Define basis for chebyshev polynomial expansion

*     Capital

arg_k(ik) = ((2*ord(ik)-1)*pi)/(2*card(ik));
x_k(ik) = cos(arg_k(ik));
cap(ik) = (lo_k+up_k+(up_k-lo_k)*x_k(ik))/2;

*     Productivity

arg_p(ip) = ((2*ord(ip)-1)*pi)/(2*card(ip));
x_p(ip) = cos(arg_p(ip));
phi(ip) = (lo_p+up_p+(up_p-lo_p)*x_p(ip))/2;

*     Define terms included in Chebyshev polynomial basis functions in the form:
```

```
*     cc(ic) * X**ce(ic)
*     where cc(ic) is the coefficient, X is the state and ce(ic), the exponent
*     for each component of Chebyshev basis functions.

$include chebyshevp_term_define

*-------------------------------------------------------------------------------
*     Present value function based on Chebyshev polynomial terms
*-------------------------------------------------------------------------------

*     Least-squares estimation
$macro PVL(kbar,phibar,ik,ip)  (sum(cp, A(cp) * \
                          sum(ic,kbar(ik,ic)$(ic.val eq cpe("cap",cp))) * \
                          sum(ic,phibar(ip,ic)$(ic.val eq cpe("phi",cp)))))

*     Value function computation
$macro PV(KCS,phitcs,ik,ip,jp) (sum(cp, A(cp) * \
                          sum(ic,KCS(ik,ip,ic)$(ic.val eq cpe("cap",cp))) * \
                          sum(ic,phitcs(jp,ic)$(ic.val eq cpe("phi",cp)))))

*     Normalized value of K used in Chebyshev Polynomials
$macro KN(ik,ip) ((K(ik,ip)-(lo_k+up_k)/2)/((up_k-lo_k)/2))


*-------------------------------------------------------------------------------
*     Transition Matrix
*-------------------------------------------------------------------------------

Table tmatrix(ip,jp)


      1           2           3           4           5
1     0.9727      0.0273      0           0           0
2     0.0041      0.9806      0.0153      0           0
3     0           0.0082      0.9837      0.0082      0
4     0           0           0.0153      0.9806      0.0041
5     0           0           0           0.0273      0.9727;


Parameters
      phit        Grid point values of approximated phi,
      phitn       Normalized grid point values of projected phi,
      phitcs      CS Polynomial terms used for value approximation;


*     stochasticity

phit(ip) = phi(ip);
phitn(ip) = ((phit(ip)-(lo_p+up_p)/2)/((up_p-lo_p)/2));

*-------------------------------------------------------------------------------
*     Apply Chebyshev polynomial algorithm on productivity
*-------------------------------------------------------------------------------

phitcs(ip,ic) = sum((pt,pst), cc(ic,pt,pst)$csp(ic,pt,pst)*
                            power(phitn(ip),ce(ic,pt,pst)$csp(ic,pt,pst)));


*-------------------------------------------------------------------------------
*     Define Bellman Equation
*-------------------------------------------------------------------------------

Variables
      OBJ               Objective
      C(ik,ip)          Consumption,
      K(ik,ip)          Subsequent period capital stock,
      U(ik,ip)          Nodal approximations of utility,
      A(cp)             Terms in the value function approximation,
      KCS(ik,ip,ic)     Chebyshev polynomial terms (ic) for capital,
```

```
134       P(ik,ip)          Shadow price of capital;

136   Equations
137       utility     Present value benefit function,
138       market      Market for current output,
139       objdef      Least squares objective,
140       k_csdef     Chebyshev polynomial terms for capital,
141       foca        First order condition for coefficient A,
142       copt        First order condition for consumption,
143       kopt        First order condition for capital,
144       udef        Defines nodal utility;

146   utility..   OBJ =e= sum((ik,ip), 1/(1-eta) * C(ik,ip)**(1-eta) +
147                         beta * sum(jp, tmatrix(ip,jp) * PV(KCS,phitcs,ik,ip,jp)));


150   market(ik,ip)..   C(ik,ip) + K(ik,ip) =e= phi(ip) * cap(ik)**alpha;

152   objdef..    OBJ =e= sum((ik,ip), sqr(PVL(kbar,phibar,ik,ip) - U(ik,ip)));

154   k_csdef(ik,ip,ic)..
155               KCS(ik,ip,ic) =e= sum((pt,pst), cc(ic,pt,pst)$csp(ic,pt,pst)*
156                                 power(KN(ik,ip),ce(ic,pt,pst)$csp(ic,pt,pst)));

158   foca(cpp).. sum((ik,ip), 2 *
159                     (PVL(kbar,phibar,ik,ip) - U(ik,ip)) *
160                           sum(ic,kbar(ik,ic)$(ic.val eq cpe("cap",cpp))) *
161                           sum(ic,phibar(ip,ic)$(ic.val eq cpe("phi",cpp)))
162               ) =e= 0;

164   copt(ik,ip)..     P(ik,ip) =g= C(ik,ip)**(-eta);

166   kopt(ik,ip)..     P(ik,ip) =g=
167                     beta/((up_k - lo_k)/2) *
168                     sum(jp, tmatrix(ip,jp) *
169                   sum(cp, A(cp) * sum(ic, phitcs(jp,ic)$(ic.val eq cpe("phi",cp))) *
170                     sum(ic,
171                       sum((pt,pst),
172                             (cc(ic,pt,pst)$(csp(ic,pt,pst)) *
173                             ce(ic,pt,pst)$(csp(ic,pt,pst)) *
174                             power(KN(ik,ip),ce(ic,pt,pst)-1))$(ce(ic,pt,pst) ge 1)
175                       )$(ic.val eq cpe("cap",cp))
176                     )
177                   )
178               );

180   udef(ik,ip)..     U(ik,ip) =e=
181                     1/(1-eta) * C(ik,ip)**(1-eta) +
182                           beta * sum(jp, tmatrix(ip,jp) * PV(KCS,phitcs,ik,ip,jp));

184   model bellman /utility,market,k_csdef/;
185   model lsqr /objdef/;
186   model oneshot_mcp /foca.A, copt.C, kopt.K, market.P, udef.U, k_csdef.KCS/;

188   *-------------------------------------------------------------------------------
189   *     Initialization
190   *-------------------------------------------------------------------------------

192   C.LO(ik,ip) = 1e-6;
193   K.LO(ik,ip) = 0;
194   A.L(cp) = 1;

196   *-------------------------------------------------------------------------------
197   *     Value Iteration
198   *-------------------------------------------------------------------------------

200   *     Parmeters for value function iteration

202   Parameters
```

```
203         dev   Current deviation coef  /1/,
204         itlog Iteration log;

206 *       Initial value:

208 U.FX(ik,ip) = 1;

210 file ktitle; ktitle.lw=0;

212 bellman.solvelink = 2;
213 loop(iter$round(dev,10),

215         itlog(iter,cp) = A.L(cp);

217         A.LO(cp) = -INF;  A.UP(cp) = +INF;

219         solve lsqr using nlp minimzing OBJ;

221         dev = sum(cp,sqr(A.L(cp)-itlog(iter,cp)));

223         itlog(iter,"dev") = dev;

225         A.FX(cp) = A.L(cp);

227         solve bellman using nlp max obj;
228         abort$(bellman.solvestat<>1 and bellman.modelstat>2) "Bellman does not solve.";
229         U.FX(ik,ip) = 1/(1-eta) * C.L(ik,ip)**(1-eta) + beta *
230                       sum(jp, tmatrix(ip,jp) *
231                             sum(cp, A.L(cp) *
232                               sum(ic, KCS.L(ik,ip,ic)$(ic.val eq cpe("cap",cp))) *
233                               sum(ic,phitcs(jp,ic)$(ic.val eq cpe("phi",cp)))
234                             )
235                       );
236         put ktitle;
237         put_utility 'title' /'Iter: ',iter.tl,' Deviation = ',dev;

239 );

241 A.LO(cp) = -inf;
242 A.UP(cp) = +inf;
243 U.UP(ik,ip) = +inf;
244 U.LO(ik,ip) = -inf;
245 P.L(ik,ip) = market.m(ik,ip);

247 solve oneshot_mcp using mcp;
248 display K.L, U.L, A.L, KCS.L;
```

2. Hydropower Planning Model

```
$title NLP–MCP Hybrid Formulation of Hydropower Planning Model:

*     GAMS code for 2nd order polynomial estimation

*     Number of reservior level nodes:

$if not set nkl $set nkl 5

*     Number of precipitation level nodes:

$if not set nkp $set nkp 3

Set    m     Months      /jan, feb, mar, apr, may, jun,
                         jul, aug, sep, oct, nov, dec /;


Parameters  inflow(m)   Mean inflow (million m3) /
                  jan         1.2,
                  feb         1.0,
                  mar         1.1,
                  apr         1.2,
                  may         40.2,
                  jun         99.5,
                  jul         146.3,
                  aug         138.2,
                  sep         70.7,
                  oct         11.7,
                  nov         2.3,
                  dec         1.5 /;

*     Set maximum capcity of reservoir

$if not set rmax $set rmax 140

Parameters
      lmax    Maximum water level the dam can store (million m3)        /250/,
      lmin    Minimum water level that must be maintained (million m3)  /0/,
      rmax    Maximum amount that can be released per month (m. m3)     /%rmax%/,
      g(m)    Monthly demand (fixed),
      cref    Reference cost /1/,
      eta         Elasticity of non–hydro supply /2/,
      xref    Reference non–hydro supply /100/;

g(m) = 140;

Set    kl    Water level grid points /0*%nkl%/;

Parameters
      L(kl,m)           Water levels at grid points,
      theta(*)          Parameter defining convex combinations;

*    Water level state variable uniformly distributed
*    between the min and max level:

theta(kl) = (ord(kl)–1)/(card(kl)–1);
L(kl,m) = lmin * (1–theta(kl)) + lmax * theta(kl);

Set         kp          Precipitation grid points /0*%nkp%/;

Parameters
      dfac              Monthly discount factor (6% per year) /0.99/
      rhod              Rainfall persistence (highly persistent) /0.9/
      sigma             Normalized standard deviation of inflow /0.1/,
      d_mean            Mean value of d
      d_std             Standard deviation of d
      d_low             Low value of d on the grid,
      d_high            High value of d on the grid,
```

```
68        d(kp,m)           Grid point values of d;


71  *     AR(1) mean:

73  d_mean(m) = inflow(m);

75  *     AR(1) std. dev.:

77  d_std(m)  = sigma * d_mean(m) /sqrt((1-rhod**2));
78  d_low(m)  = d_mean(m) - 3*d_std(m);
79  d_high(m) = d_mean(m) + 3*d_std(m);

81  *     Set up the Gauss-Hermite grid:

83  Set   i     Gaussian-Hermite grid points /1*5/


86  *     Implementing GH:
87  *     1. Get n nodes and n weights from a computer program

89  Table ghdata(i,*)    Tabulated Gaussian-Hermite approximation points

91          zeta    omega
92     1    2.0202  0.02
93     2    0.9586  0.3936
94     3    0         0.9453
95     4    -0.9586 0.3936
96     5    -2.0202 0.02;

98  Parameter
99        omega(i)          Normalized GH weights,
100       zeta(m,i)         Normalized GH nodes,
101       dt(kp,m,i)        Grid point values of dt;

103 *     Take expectation of function value:
104 *     omega(i) * h(zeta)

106 omega(i) = ghdata(i,"omega")/sqrt(3.141592);
107 zeta(m,i) = sigma*sqrt(2)*ghdata(i,"zeta");

109 *     discretize precipitation levels

111 theta(kp) = (ord(kp)-1)/(card(kp)-1);
112 d(kp,m) = d_low(m) * (1-theta(kp)) + d_high(m) * theta(kp);

114 *     Rainfall in realization i following month m is
115 *     defined relative to the mean rainfall in month m++1.
116 *     zeta is normally distributed:

118 dt(kp,m++1,i) = max(0, d_mean(m++1) * (1 + rhod * (d(kp,m)/d_mean(m)-1) + zeta(m++1,i)));

120 Set   ja /0*2/    Set index for coefficients: L,
121       jb /0*2/    Set index for coefficients: d,
122       k(kl,kp)    State of the world;

124 k(kl,kp) = yes;

126 alias (ja,ja_), (jb,jb_);

128 *     First order taylor linear approximation of the optimal
129 *     present value of month m with reservoir level L and precipitation d:
130 *     searching for taylor approximation coefficients

132 $macro PV(m,L,d) (sum(ja_,A(ja_,m)*power(L,ja_.val))+sum(jb_,B(jb_,m)*power(d,jb_.val)))

134 Variables
135       P(kl,kp,m)        Current estimate of the shadow price of water,
136       A(ja,m)           Polynomial approximation terms for L,
```

```
137         B(jb,m)             Polynomial approximation terms for d,
138         PE(kl,kp,m)         Shadow price of electricity,
139         X(kl,kp,m)          Non-hydro generation,
140         Z(kl,kp,m)          Water retained,
141         R(kl,kp,m)          Water released to generate electricity,
142         S(kl,kp,m,i)        Water spilled without generating electricity,
143         PT(kl,kp,m,i)       Projected shadow price,
144         MU(kl,kp,m,i)       Shadow price on upper bound constraint,
145         LT(kl,kp,m,i)       Projected level;

147 Free Variable
148         OBJ          NLP Objective;

150 Nonnegative Variables Z,R,S;

152 Equations    objdef, supply, demand, ptdef, supplyt, upper, lsqrdef;

154 *       NLP Objective Definition

156 objdef..            OBJ =e= sum((k(kl,kp),m),

158 *       Discounted expected value of subsequent state:

160                     dfac * sum(i, omega(i) *

162 *       Value of water delivered in the subsequent month:

164                     PT(k,m++1,i) * LT(k,m++1,i))

166 *       Cost of using non-hydro generation this month:

168                     - cref * xref * (X(k,m)/xref)**(eta) );


171 *       Water at the start of the month is either retained or released to
172 *       generate electricity:

174 supply(k(kl,kp),m)..     L(kl,m) =g= Z(k,m) + R(k,m);

176 *       Total generation equals that provided by hydro and non-hydro
177 *       sources.  Demand is fixed in each month:

179 demand(k(kl,kp),m)..     g(m) =e= R(k,m) + X(k,m);

181 *       The level of water at the start of month m depends on how much water
182 *       was stored the previous month (Z), how much inflow occurred (dt) and
183 *       how much water was spilled (S):

185 supplyt(k(kl,kp),m,i).. LT(k,m,i) =e= Z(k,m--1) + dt(kp,m,i) - S(k,m,i);

187 *       The price of water in the subsequent month is imputed on the basis of
188 *       the imputed water price on the nodes along with the shadow prices of
189 *       the upper and lower bounds on capacity:

191 ptdef(k(kl,kp),m,i)..   PT(k,m,i) =e= PV(m, LT(k,m,i), dt(kp,m,i));

193 *       Least-squares estimation of shadow price of water

195 lsqrdef..               OBJ =E= sum((k(kl,kp),m), sqr(PV(m,L(kl,m),d(kp,m)) - p(k,m)));

197 Equations   foc_a(ja,m) First order condition for a (MCP formlation)
198             foc_b(jb,m) First order condition for b (MCP formulation);

200 foc_a(ja,m)..     -sum(k(kl,kp),(P(k,m)-PV(m,L(kl,m),d(kp,m))) * power(L(kl,m),ja.val)) =e= 0;

202 foc_b(jb,m)..     -sum(k(kl,kp),(P(k,m)-PV(m,L(kl,m),d(kp,m))) * power(d(kp,m),jb.val)) =e= 0;

204 *       Dual equations for MCP
```

```
206 Equations   xopt, ropt, zopt, sopt;

208 *     Marginal cost of non-hydro generation equals market price of
209 *     electricity:

211 xopt(k(kl,kp),m)..      cref * eta * (X(k,m)/xref)**(eta-1) =g= PE(k,m);

213 *     Hydro generation is equalized with the shadow value of water:

215 ropt(k(kl,kp),m)..      P(k,m) =g= PE(k,m);

217 *     Risk neutral: water left in the reservior equals the expected
218 *     value of water in the subsequent month:

220 zopt(k(kl,kp),m)..      P(k,m) =g= dfac * sum(i, omega(i) *
221                                          (PT(k,m++1,i) +
222                                                LT(k,m++1,i)*(A("1",m++1) +
223                                                      2*A("2",m++1)*LT(k,m++1,i)) ));

225 *     Spilling water amounts to free disposal.  This assures that
226 *     the shadow value of water in subsequent months is nonnegative:

228 sopt(k(kl,kp),m,i)..    dfac * omega(i) * (PT(k,m,i) +
229                                    LT(k,m,i)*(A("1",m) +
230                                          2*A("2",m)*LT(k,m,i)) ) =g= 0;


233 model       hydronlp    /objdef,supply,demand,supplyt,ptdef/;
234 model       lsqr        /lsqrdef/;
235 model       oneshotmcp  /foc_a.A,foc_b.B,supply.P,demand.PE,
236                          supplyt.LT,ptdef.PT,sopt.S,xopt.X,
237                          zopt.Z,ropt.R/;

239 *    Initialial values for a "cold" start:

241 A.L("0",m) = cref*eta;
242 B.L(jb,m) = 0;
243 P.L(kl,kp,m) = cref*eta;
244 PE.L(kl,kp,m) = cref*eta;
245 PT.L(kl,kp,m,i) = cref*eta;
246 LT.L(kl,kp,m,i) = 0.5*(lmax-lmin);
247 R.L(kl,kp,m) = rmax;
248 X.L(kl,kp,m) = xref;
249 Z.L(kl,kp,m) = 0.4*(lmax-lmin);
250 X.LO(kl,kp,m) = 0.001;
251 LT.UP(k,m,i) = lmax;

253 *     Initial estimate at value of water:

255 P.FX(kl,kp,m) = cref*eta;


258 set         iter /1*30/;

260 alias (kl,kl_), (kp,kp_), (m,m_);

262 file ktitle; ktitle.lw=0;

264 parameter   pivotdata(*,kl,kp,m), coef(*,*,m);

266 parameter         itlog    Iteration log;

268 loop(iter,
269    itlog(kl,kp,m,iter) = P.L(kl,kp,m);

271    put ktitle;
272    put_utility  'title' / 'Solving for value function coefficients, iteration ',iter.tl;
273    A.LO(ja,m) = -INF;    A.UP(ja,m) = +INF;
274    B.LO(jb,m) = -INF;    B.UP(jb,m) = +INF;
```

```
276     solve lsqr using nlp minimzing OBJ;

278     A.FX(ja,m) = A.L(ja,m);     B.FX(jb,m) = B.L(jb,m);

280     put_utility 'title' / 'Solving dynamic programming recursion.';

282     option solprint = on;
283     hydronlp.savepoint = 1;
284     solve hydronlp max OBJ using nlp;
285     abort$(hydronlp.modelstat>2) "Model fails to solve.";

287     P.FX(kl,kp,m) = -supply.m(kl,kp,m);

289 );

291 *     Free up estimated coefficients and price value

293 A.LO(ja,m) = -INF;       A.UP(ja,m) = +INF;
294 B.LO(jb,m) = -INF;       B.UP(jb,m) = +INF;
295 P.LO(kl,kp,m) = -INF;    P.UP(kl,kp,m) = +INF;

297 *     Read in dual variables levels from NLP

299 P.L(kl,kp,m) = -supply.m(kl,kp,m);
300 PE.L(kl,kp,m) = -demand.m(kl,kp,m);
301 LT.L(kl,kp,m,i) = supplyt.m(kl,kp,m,i);
302 PT.L(kl,kp,m,i) = ptdef.m(kl,kp,m,i);

304 *     Solve MCP

306 solve oneshotmcp using mcp;
```

Chapter 6.

Conclusions

This dissertation introduces a stochastic control framework (DICESC) to investigate the effect of the risk and uncertainty of climate thresholds on policy. Stochastic control, is in essence, defined by the its welfare maximizing agent's ability to *learn*. In chapter 2, I show that the model components that embody learning—namely, endogenous hazard rates (active learning) in an act-then learn framework (passive learning)—is a powerful driver of precautionary abatement. More importantly in chapter 3, I demonstrate that learning in stochastic control is actually fundamental in modeling how we make decisions under the risk of tipping. Implementations of the DICE model in the absence of either type of learning prove that the optimal policies output by lack-of-learning frameworks can lead to determining the wrong sign, if not the wrong magnitude of the uncertainty effect.

Despite the fundamental significance of applying stochastic control however, model results in chapter 2 indicate that the approach suffers from the pitfalls in Bayesian learning—even in the most minimal frameworks. Although the incentive for precautionary measures is robust, distributional sensitivity analyses expose the stochastic framework's strong dependency on higher moments of the Bayesian prior; i.e., two distributions showing the same mean and standard deviation may argue for very different policy recommendations. A decomposed analysis of abatement reveals that precautionary abatement incentives are sensitive to higher moment distributive properties. The dependency on higher moments is more

severe when assessing the expected value of information. Humility is called for as it seems unlikely that we can determine the distribution of risk with sufficient precision.

Insufficient precision in the Bayesian prior has naturally lead the IAM community to focus on the notion of (Knightian) uncertainty, or ambiguity (chapter 4). Given multiple distributions of risk, I implement four models of ambiguity aversion in DICESC and derive optimal carbon abatement policies for each ambiguity attitude. Not surprisingly, results indicate a robust precautionary incentive to heighten policy under ambiguity aversion. However, in a multi-prior, stochastic control setting, there exist two types of priors—time-specific probabilities of tipping and state-specific probabilities assigned to each state of the world. I find that hedging strategies can differ significantly depending on which type of prior ambiguous information is established with. More importantly, because priors are assumed to be convex combinations of the prior set in each period, an abrupt transition between priors can impose threshold effects on the optimal policy even before tipping occurs; an ambiguity averse decision maker exhibits incentives to delay reaching temperature points at which the assumed Bayesian prior takes a turn for the worse.

In chapter 5, I introduce a mixed complementarity formulation of dynamic programming problems. The primary objective of this chapter is pedagogic; to provide a technically creative alternative to the traditional value function iteration algorithm used widely in the climate IAM literature. Instead of iterating the NLP based diagonalization procedure employed in value function iteration, DP-MCP solves a system of equilibrium constraints that characterize a Nash equilibrium in the two subproblem types of the value iteration procedure. These subproblems are, namely *a.* solving for Bellman's principle of optimality, conditional on the estimate of the value function; and *b.* optimizing the parameters of the value function estimate given the vector of optimized state variables.

By omitting the iterative diagonalization process that drives conventional DP formulations, the oneshot DP-MCP approach greatly reduces the run-time required to solve DP problems. I demonstrate the computational advantage of the MCP approach by extending the application to a computationally burdensome stochastic growth problem. The MCP approach also allows for the proper treatment of corner solutions via complementary slackness conditions when solving for the optimal policy, making the complementarity formulation both efficient and robust.

References

Extensions of GAMS for complementarity problems arising in applied economics. *Journal of Economic Dynamics and Control*, 19(8):1299–1324, 1995.

Frank Ackerman, Elizabeth A Stanton, and Ramón Bueno. Fat tails, Exponents, Extreme Uncertainty: Simulating Catastrophe in DICE. *Ecological Economics*, 69(8):1657–1665, 2010.

Frank Ackerman, Elizabeth A Stanton, and Ramón Bueno. Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy? *Environmental and Resource Economics*, 56(1): 73–84, 2013.

Nabil I Al-Najjar and Jonathan Weinstein. The Ambiguity Aversion Literature: A Critical Assessment. *Economics and Philosophy*, 25(03):249–284, 2009.

Bita Analui and Georg Ch Pflug. On Distributionally Robust Multiperiod Stochastic Optimization. *Computational Management Science*, 11(3):197–220, 2014.

Kenneth J Arrow and Anthony C Fisher. Environmental Preservation, Uncertainty, and Irreversibility. In *Classic Papers in Natural Resource Economics*, pages 76–84. Springer, 1974.

S Borağan Aruoba and Jesús Fernández-Villaverde. A comparison of programming languages in economics. Technical report, National Bureau of Economic Research, 2014.

Erin Baker. Optimal Policy Under Uncertainty and Learning About Climate Change: A Stochastic Dominance Approach. *Journal of Public Economic Theory*, 11(5):721–747, 2009.

Edward J Balistreri. Playing the bertrand game in the corner: A mixed complementarity formulation with endogenous entry. 1999.

Ravi Bansal and Amir Yaron. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

William J Baumol. Activity analysis in one lesson. *The American Economic Review*, 48(5):837–873, 1958.

R Bellman. Dynamic programming. rand corporation research study. *Princeton University Press, 1957b. ISBN*, 1050546524:3, 1957.

Loïc Berger, Johannes Emmerling, and Massimo Tavoni. Managing Catastrophic Climate Risks Under Model Uncertainty Aversion. *Management Science*, 2016.

Valentina Bosetti, Emanuele Massetti, and Massimo Tavoni. The WITCH model. Structure, Baseline, Solutions. 2007.

Paolo Brandimarte. *Handbook in Monte Carlo Simulation: Applications in Financial Engineering, Risk Management, and Economics*. John Wiley & Sons, 2014.

William A Brock and Leonard J Mirman. Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory*, 4(3):479–513, 1972.

Yongyang Cai and Kenneth L Judd. Shape-preserving dynamic programming. *Mathematical Methods of Operations Research*, 77(3):407–421, 2013.

Yongyang Cai and Kenneth L Judd. Dynamic programming with hermite approximation. *Mathematical Methods of Operations Research*, 81(3):245–267, 2015.

Yongyang Cai, Kenneth L Judd, and Thomas S Lontzek. DSICE: A Dynamic Stochastic Integrated Model of Climate and Economy. *Center for Robust Decision-Making on Climate Policy, Working Paper*, (12-06), 2012.

Wonjun Chang and Thomas F Rutherford. Catastrophic Thresholds, Bayesian Learning and the Robustness of Climate Policy Recommendations. 2017.

Graciela Chichilnisky and Geoffrey Heal. Global Environmental Risks. In *Sustainability: Dynamics and Uncertainty*, pages 23–46. Springer, 1998.

Benjamin Crost and Christian P Traeger. Risk and Aversion in the Integrated Assessment of Climate Change. *Department of Agricultural & Resource Economics, UCB*, 2011.

Benjamin Crost and Christian P Traeger. Optimal Climate Policy: Uncertainty Versus Monte Carlo. *Economics Letters*, 120(3):552–558, 2013.

Delavane Diaz and Klaus Keller. A potential disintegration of the west antarctic ice sheet: Implications for economic analyses of climate policy. *The American Economic Review*, 106(5): 607–611, 2016.

Simon Dietz. High Impact, Low Probability? An Empirical Analysis of Risk in the Economics of Climate Change. *Climatic Change*, 108(3):519–541, 2011.

Simon Dietz and Nicholas Stern. Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus' Framework Supports Deep Cuts in Carbon Emissions. *The Economic Journal*, 125(583):574–620, 2015.

Steven P Dirkse and Michael C Ferris. The PATH solver: a nommonotone stabilization scheme for mixed complementarity problems. *Optimization Methods and Software*, 5(2):123–156, 1995.

Laurent Drouet, Valentina Bosetti, and Massimo Tavoni. Selection of Climate Policies Under the Uncertainties in the Fifth Assessment Report of the IPCC. *Nature climate change*, 5(10): 937–940, 2015.

Jean-Pierre Dubé, Jeremy T Fox, and Che-Lin Su. Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica*, 80(5):2231–2267, 2012.

Daniel Ellsberg. Risk, Ambiguity, and the Savage Axioms. *The quarterly journal of economics*, pages 643–669, 1961.

Larry G Epstein. Decision Making and the Temporal Resolution of Uncertainty. *International economic review*, pages 269–283, 1980.

Larry G Epstein and Martin Schneider. Recursive Multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.

Larry G Epstein and Stanley E Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica: Journal of the Econometric Society*, pages 937–969, 1989.

Ronald Fagin and Joseph Y. Halpern. Uncertainty, Belief, and Probability. *Computational Intelligence*, 7(3):160–173, 1991.

Michael C Ferris and Todd S Munson. Complementarity problems in gams and the PATH solver. *Journal of Economic Dynamics and Control*, 24(2):165–188, 2000.

Michael C Ferris and Krung Sinapiromsaran. Formulating and solving nonlinear programs as mixed complementarity problems. In *Optimization*, pages 132–148. Springer, 2000.

Michael C Ferris, Steven P Dirkse, Jan-H Jagla, and Alexander Meeraus. An extended mathematical programming framework. *Computers & Chemical Engineering*, 33(12):1973–1982, 2009.

Itzhak Gilboa and David Schmeidler. Maxmin Expected Utility with Non-unique Prior. *Journal of mathematical economics*, 18(2):141–153, 1989.

Lars Peter Hansen and Thomas J Sargent. *Robustness*. Princeton university press, 2008.

Geoffrey Heal and Antony Millner. Uncertainty and Decision in Climate Change Economics. Technical report, National Bureau of Economic Research, 2013.

Geoffrey Heal and Antony Millner. Uncertainty and decision making in climate change economics. *Review of Environmental Economics and Policy*, 8(1):120–137, 2014.

Tito Homem-de Mello and Güzin Bayraksan. Monte Carlo Sampling-based Methods for sTochastic Optimization. *Surveys in Operations Research and Management Science*, 19(1):56–85, 2014.

Chris Hope. The Marginal Impact of CO2 from PAGE2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern. *Integrated assessment*, 6(1):19–56, 2006.

Richard Howitt, Siwa Msangi, Arnaud Reynaud, and Keith Knapp. Using polynomial approximations to solve stochastic dynamic programming problems: Or a 'betty crocker' approach to sdp. *University of California, Davis*, 2002a.

Richard E Howitt, Arnaud Reynaud, Siwa Msangi, Keith C Knapp, et al. Calibrated stochastic dynamic models for resource management. In *The 2nd World Congress of Environmental and Resource Economists*, volume 2427, 2002b.

Terrence Iverson. Optimal Carbon Taxes with Non-constant Time Preference. 2012.

Svenn Jensen and Christian Traeger. Growth and Uncertainty in the Integrated Assessment of Climate Change. *Eur. Econom. Rev.(in the press, 2014)*, 2011.

Kenneth L Judd. *Numerical methods in economics*. MIT press, 1998.

Kenneth L Judd, Lilia Maliar, Serguei Maliar, and Rafael Valero. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. *Journal of Economic Dynamics and Control*, 44:92–123, 2014.

Noah Kaufman. The Bias of Integrated Assessment Models that Ignore Climate Catastrophes. *Climatic change*, 110(3-4):575–595, 2012.

Klaus Keller, Benjamin M Bolker, and David F Bradford. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management*, 48(1):723–741, 2004.

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

Anton J Kleywegt, Alexander Shapiro, and Tito Homem-de Mello. The Sample Average Approximation Method for Stochastic Discrete Optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A Smooth Model of Decision Making Under Ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive Smooth Ambiguity Preferences. *Journal of Economic Theory*, 144(3):930–976, 2009.

Charles D Kolstad. *Learning, Irreversibilities and Climate*. Citeseer, 1996.

Karen A Kopecky. Function approximation. *University of Western Ontario Lectures Notes ECO 613/614 Fall 2007*, 2007.

Elmar Kriegler, Jim W Hall, Hermann Held, Richard Dawson, and Hans Joachim Schellnhuber. Imprecise Probability Assessment of Tipping Points in the Climate System. *Proceedings of the National Academy of Sciences*, 106(13):5041–5046, 2009.

Andreas Lange and Nicolas Treich. Uncertainty, Learning and Ambiguity in Economic Models on Climate Policy: Some Classical Results and New Directions. *Climatic Change*, 89(1): 7–21, 2008.

Morten I Lau, Andreas Pahlke, and Thomas F Rutherford. Approximating Infinite-horizon Models in a Complementarity Format: A Primer in Dynamic General Equilibrium Analysis. *Journal of Economic Dynamics and Control*, 26(4):577–609, 2002.

Andrew J Leach. The Climate Change Learning Curve. *Journal of Economic Dynamics and Control*, 31(5):1728–1752, 2007.

Derek Lemoine and Ivan Rudik. Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. 2016.

Derek Lemoine and Christian Traeger. Watch Your Step: Optimal Policy in a Tipping Climate. *American Economic Journal: Economic Policy*, 6(1):137–166, 2014.

Derek Lemoine and Christian P Traeger. Ambiguous Tipping Points. *Journal of Economic Behavior & Organization*, 132:5–18, 2016a.

Derek Lemoine and Christian P Traeger. Economics of Tipping the Climate Dominoes. *Nature Climate Change*, 2016b.

Derek M Lemoine and Christian Traeger. Tipping Points and Ambiguity in the Integrated Assessment of Climate Change. *University of California, Berkeley, Department of Agricultural and Resource Economics working paper*, 1111, 2010.

Timothy M Lenton. Early Warning of Climate Tipping Points. *Nature Climate Change*, 1(4): 201–209, 2011.

Timothy M Lenton. Arctic Climate Tipping Points. *Ambio*, 41(1):10–22, 2012.

Timothy M Lenton, Hermann Held, Elmar Kriegler, Jim W Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber. Tipping Elements in the Earth's Climate System. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, 2008.

Thomas S Lontzek, Yongyang Cai, and Kenneth L Judd. Tipping Points in a Dynamic Stochastic IAM. 2012.

Mark J Machina and Marciano Siniscalchi. Ambiguity and Ambiguity Aversion. *Handbook of the Economics of Risk and Uncertainty*, 1:729–807, 2014.

Lilia Maliar and Serguei Maliar. Merging simulation and projection approaches to solve high-dimensional problems with an application to a new keynesian model. *Quantitative Economics*, 6(1):1–47, 2015.

Alan Manne, Robert Mendelsohn, and Richard Richels. MERGE: A Model for Evaluating Regional and Global Effects of GHG Reduction Policies. *Energy policy*, 23(1):17–34, 1995.

Alan S Manne and Richard G Richels. MERGE: An Integrated Assessment Model for Global Climate Change. In *Energy and Environment*, pages 175–189. Springer, 2005.

Alan Sussmann Manne and Richard G Richels. *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emission Limits*. MIT Press, 1992.

Rodolfo E Manuelli and Thomas J Sargent. *Exercises in dynamic macroeconomic theory*. Harvard University Press, 2009.

Ian WR Martin and Robert S Pindyck. Averting Catastrophes: The Strange Economics of Scylla and Charybdis. *The American Economic Review*, 105(10):2947–2985, 2015.

John C Mason and David C Handscomb. *Chebyshev polynomials*. CRC Press, 2002.

Malte Meinshausen, Nicolai Meinshausen, William Hare, Sarah CB Raper, Katja Frieler, Reto Knutti, David J Frame, and Myles R Allen. Greenhouse-gas Emission Targets for Limiting Global Warming to 2 C. *Nature*, 458(7242):1158–1162, 2009.

Antony Millner, Simon Dietz, and Geoffrey Heal. Scientific Ambiguity and Climate Policy. *Environmental and Resource Economics*, 55(1):21–46, 2013.

Carmen G Moles, Julio R Banga, and Klaus Keller. Solving Nonconvex Climate Control Problems: Pitfalls and Algorithm Performances. *Applied Soft Computing*, 5(1):35–44, 2004.

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

W Nordhaus and P Sztorc. DICE 2013: Introduction and User's Manual, 2014.

William D Nordhaus. *Managing the Global Commons: The Economics of Climate Change*. MIT press Cambridge, MA, 1994.

William D Nordhaus. The Economic Impacts of Abrupt Climatic Change. *Prepared for Meeting on Abrupt Climate Change: The Role of Oceans, Atmosphere, and the Polar Regions, National Research Council*, 1999.

William D Nordhaus. Integrated Economic and Climate Modeling. Technical report, Cowles Foundation for Research in Economics, Yale University, 2011.

William D Nordhaus. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, 2014.

Beresford N Parlett, H Simon, and LM Stringer. On estimating the largest eigenvalue with the lanczos algorithm. *Mathematics of computation*, 38(157):153–165, 1982.

Robert S Pindyck. Climate Change Policy: What do the Models Tell Us? *Journal of Economic Literature*, 51(3):860–872, 2013.

Warren B Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, volume 703. John Wiley & Sons, 2007.

Warren B Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, volume 842. John Wiley & Sons, 2011.

Michael Rothschild and Joseph E Stiglitz. Increasing Risk: I. A Definition. *Journal of Economic Theory*, 2(3):225–243, 1970.

Ivan Rudik. The social cost of carbon when damages are unknown. Technical report, Working paper, Iowa State University, Ames, 2016.

John Rust. Numerical dynamic programming in economics. *Handbook of computational economics*, 1:619–729, 1996.

Andrzej Ruszczynski and Alexander Shapiro. Conditional Risk Mappings. *Mathematics of operations research*, 31(3):544–561, 2006.

Yousef Saad. On the rates of convergence of the lanczos and the block-lanczos methods. *SIAM Journal on Numerical Analysis*, 17(5):687–706, 1980.

Thomas Sargent and John Stachurski. Quantitative economics with python. Technical report, Technical report, Lecture Notes, 2015.

J Savage Leonard. The Foundations of Statistics. *NY, John Wiley*, pages 188–190, 1954.

Herbert Scarf, KJ Arrow, and S Karlin. A Min-max Solution of an Inventory Problem. *Studies in the mathematical theory of inventory and production*, 10:201–209, 1958.

Herbert Scarf et al. *On the computation of equilibrium prices*. Number 232. Cowles Foundation for Research in Economics at Yale University, 1967.

David Schmeidler. Subjective Probability and Expected Utility without Additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587, 1989.

Stephen H Schneider. *Climate Change Science and Policy*. Island Press, 2009.

Alexander Shapiro. On a Time Consistency Concept in Risk Averse Multistage Stochastic Programming. *Operations Research Letters*, 37(3):143–147, 2009.

Alexander Shapiro and Anton Kleywegt. Minimax Analysis of Stochastic Problems. *Optimization Methods and Software*, 17(3):523–542, 2002.

SM Sinha. *Mathematical Programming: Theory and Methods*. Elsevier, 2005.

Joel B Smith, Stephen H Schneider, Michael Oppenheimer, Gary W Yohe, William Hare, Michael D Mastrandrea, Anand Patwardhan, Ian Burton, Jan Corfee-Morlot, Chris HD Magadza, et al. Assessing Dangerous Climate Change Through an Update of the Intergovernmental Panel on Climate Change (IPCC) "Reasons for Concern". *Proceedings of the National Academy of Sciences*, 106(11):4133–4137, 2009.

Thomas F Stocker, Dahe Qin, Gian-Kasper Plattner, Melinda Tignor, Simon K Allen, Judith Boschung, Alexander Nauels, Yu Xia, Vincent Bex, Pauline M Midgley, et al. 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. Technical report, 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), 2013.

Nancy L Stokey and Robert E with Edward C. Prescott Lucas Jr. Recursive methods in economic dynamics, 1989.

Robert Henry Strotz. Myopia and Inconsistency in Dynamic Utility Maximization. *The Review of Economic Studies*, pages 165–180, 1955.

Che-Lin Su and Kenneth L Judd. Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5):2213–2230, 2012.

George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181, 1986.

Richard SJ Tol. On the Optimal Control of Carbon Dioxide Emissions: An Application of FUND. *Environmental Modeling & Assessment*, 2(3):151–163, 1997.

Richard SJ Tol. The Economic Effects of Climate Change. *The Journal of Economic Perspectives*, pages 29–51, 2009.

Christian P Traeger. Subjective Risk, Confidence, and Ambiguity. 2011.

Nathan M Urban, Philip B Holden, Neil R Edwards, Ryan L Sriver, and Klaus Keller. Historical and Future Learning About Climate Sensitivity. *Geophysical Research Letters*, 41(7): 2543–2552, 2014.

Annette Vissing-Jørgensen and Orazio P Attanasio. Stock-market Participation, Intertemporal Substitution, and Risk-aversion. *American Economic Review*, pages 383–391, 2003.

Mort Webster. The Curious Role of "Learning" in Climate Policy: Should We Wait for More Data? *The Energy Journal*, pages 97–119, 2002.

Mort Webster, Nidhi Santen, and Panos Parpas. An Approximate Dynamic Programming Framework for Modeling Global Climate Policy Under Decision-dependent Uncertainty. *Computational Management Science*, 9(3):339–362, 2012.

Martin L Weitzman. On Modeling and Interpreting the Economics of Catastrophic Climate Change. *The Review of Economics and Statistics*, 91(1):1–19, 2009.

Martin L Weitzman. Fat-tailed Uncertainty in the Economics of Catastrophic Climate Change. *Review of Environmental Economics and Policy*, 5(2):275–292, 2011.

Stephen Wright and Jorge Nocedal. Numerical optimization. *Springer Science*, 35:67–68, 1999.

Kirsten Zickfeld, Anders Levermann, M Granger Morgan, Till Kuhlbrodt, Stefan Rahmstorf, and David W Keith. Expert Judgements on the Response of the Atlantic Meridional Overturning Circulation to Climate Change. *Climatic Change*, 82(3-4):235–265, 2007.