The Social Cost Of Abrupt Climate Change *

Yongyang Cai[†] Kenneth L. Judd[‡] Thomas S. Lontzek[§]

30th August 2012

Abstract

We analyze a standard dynamic stochastic growth model with climate uncertainty. More specifically, we generalize DSICE, a DSGE full-dimensional extension of the DICE2007 model of William Nordhaus, by adding climate tipping events of uncertain impact on productivity and allowing Epstein-Zin preferences. We find that empirically plausible Epstein-Zin preferences imply significantly higher carbon taxes than implied by DICE2007. Moreover, the carbon tax is very sensitive to uncertainty in damage from a climate tipping event when we use standard Epstein-Zin parameters. Generally, using empirically plausible Epstein-Zin preferences implies substantially higher carbon taxes than DICE2007 and models with similar specifications. We find that these results hold when we add business-cycle shocks and multidimensional tipping processes. DSICE is solved with stable and fast dynamic programming methods with a one-year time period, with six or seven continuous state variables and several discrete state variables. This analysis indicates that there is much greater urgency to immediately enacting significant GHG policies than implied by DICE2007 and similar models that ignore uncertainty.

JEL Classification: C63, Q54, D81

Keywords: Climate policy, social costs of carbon, carbon tax, stochastic IAM, abrupt and irreversible climate change, multiple tipping points, Epstein-Zin preferences, business cycle

*We want to thank Kenneth Arrow, Paul David and Larry Goulder for comments on earlier versions of the paper. We are also grateful for comments from audiences at the AEA (2011 Denver) conference and the Berkeley ERE, and Stanford SEEPAC seminars. Furthermore, Cai, Judd, and Lontzek gratefully acknowledge NSF support (SES-0951576). Financial support for Lontzek was provided by the Zürcher Universitätsverein and the Ecosciencia Foundation. Part of this study was done while Lontzek was visiting the Hoover Institution.

[†]Hoover Institution & NBER, 424 Galvez Mall, Stanford University, Stanford, CA 94305. yycai@stanford.edu

[‡]Hoover Institution & NBER, 24 Galvez Mall, Stanford University, Stanford, CA 94305. kennethjudd@mac.com

[§]University of Zurich, Moussonstrasse 15, 8044 Zurich. thomas.lontzek@business.uzh.ch

1 Introduction

The global climate system is complex and poorly understood, particularly in terms of its response to future increases in anthropogenic GHGs. Policy choices will be affected by the uncertainty in the damage of global warming on economic productivity and when such damage is experienced. One possible policy is a carbon tax, which is a convenient one to examine since the optimal carbon tax equals the social marginal cost of the marginal ton of carbon. In 2010 the U.S. Government Interagency Working Group on Social Cost of Carbon (IWG, 2010) released an analysis of the social cost of carbon (SCC), and concluded that the 2010 SCC lies in the range \$5 - \$65 for the year 2010, with central price of \$21. Only recently, the Australian government has introduced a carbon tax of \$23. These examples demonstrate an increasing awareness of policy makers about climate policy and carbon pricing.

The IWG report relied on SCC studies in the climate and economics literature, with Nordhaus (2008) and Anthoff et al. (2009) being particularly important for the IWG. However, the IWG report, along with the majority of the literature, assumed that the economy and climate systems evolve deterministically. This ignores the substantial uncertainty that all decision makers face.

No study of the SCC models the stochastic nature of the climate and economic system in the manner and precision typically used in modern macroeconomics. Such an analysis would necessarily combine an advanced formulation of risk preferences consistent with empirical evidence with decision making processes at a time scale compatible with real world decision making. This is necessary for any realistic assessment of the costs of anthropogenic perturbations of current and future climate.

One dimension of realism is the intrinsically uncertain nature of future states of the economy. Analyses of any climate policy should account for its possible impacts on economic decision makers who face economic risks. Many issues, such as how GHG policies should react to economic fluctuations, can be studied only in models with time periods of one year or less, not in models where the time step is measured in decades.

Another dimension includes the uncertainty about the future evolution of the climate system. Several prominent studies (e.g., Nordhaus 2008) assume that damages are a function of contemporaneous temperature. However, possible climate change externalities are more complex. Many scientists are worried about climate change triggering abrupt and irreversible events leading to significant and long-lasting damages (see, e.g., Kriegler et al. 2009). Some elements of the climate system which might exhibit such a triggering effect are called tipping elements. In a prominent study, Lenton et al. (2008) characterizes tipping points for some major elements of the climate system. Examples of some tipping elements are: the weakening or shut down of the North Atlantic thermohaline circulation (THC), the melting of the Greenland ice sheet (GIS) and the West Antarctic ice sheet (WAIS), the die-back of the Amazon rainforest (AMAZ) and the increasing frequency and amplitude of El Niño-Southern Oscillation (ENSO).

While the likelihood of a tipping point occurring may be a function of contemporaneous temperature, their effects are long lasting and might be independent of future temperatures. It is assumed that some of these tipping points might occur even in this century and some tipping elements might exhibit domino dynamics and increase the tipping probability of other tipping elements (see Lenton, 2011). A faithful representation of the possibility of tipping points for the calculation of SCC would require (i) a fully stochastic multidimensional formulation of abrupt change, and (ii) modeling of the irreversibility and mutual dependencies of climate catastrophes.

The third component towards a more realistic assessment of the SCC is the modeling of the cost of risk in a manner more compatible with empirical evidence about social risk preferences. Since riskiness is inherent in the nature of tipping points, the socio-economic effects of stochastic abrupt and irreversible climate change will be affected by preferences about risks. We know from the equity premium literature that the standard formulations of preferences might misspecify how people feel about risk. Kreps and Porteus (1978) have argued that there could be value in early resolution of uncertainty, and Epstein-Zin (1989) preferences have explored the implications of this for asset pricing.

We account for all three components of a coherent analysis of climate policy under intrinsic uncertainty. This study builds on Cai, Judd and Lontzek (2012b) which combines standard features of DSGE models - productivity shocks, dynamic optimizing agents, and short time periods - with DICE2007, a seminal and basic Integrated Assessment Model (IAM). This study uses the computational framework of Cai, Judd and Lontzek (2012b) to address basic questions about uncertainty and the social cost of carbon. In particular, we present a dynamic stochastic model which incorporates the mutual interplay between climate and economics. We model a stochastic business cycle at the annual time scale, a system of stochastic tipping points, and use Epstein-Zin preferences to model social preferences.

In general, IAMs study the interplay between the climate and the economic system. All large-scale IAMs have a complex representation of the climate system but many consider economic activity as exogenously given by some pre-defined scenarios. Examples are MAGICC (Wigley and Raper, 1997), ICAM (Dowlatabati and Morgan (1993) and IMAGE (Batjes and Goldewijk, 1994). These models neglect endogenous and forward-looking decisions and are unable to account for economic reactions to climate policies and cannot study dynamic decision-making in an evolving and uncertain world. On the contrary, only a few models rely on solving an intertemporal optimization problem assuming that climate and the economy are endogenous and mutually dependent. Examples are DICE (Nordhaus, 2008), MERGE (Manne and Richels 2005) and RICE (Nordhaus and Yang, 1996). It is the latter class of models which is suitable for an uncertainty analysis in the climate policy decision making process.

Uncertainty analysis is regarded as an elementary part of these IAMs. Unfortunately, uncertainty analysis is mainly restricted to parametric uncertainty. Examples are Nordhaus (1994), Nordhaus (2008) and Pizer (1999). For general reviews, see

Heal and Kriström (2002) and Pindyck (2007). Under parametric uncertainty, the modeler doesn't know the value of key parameters, e.g., climate sensitivity or some damage function parameter. The modeler has an idea about the distribution of these parameters and conducts a Monte Carlo analysis with many simulations, each being itself a deterministic run of the model with a picked set of parameter values. This method can provide valuable information about a possible range of the climate-economy system, but the models assume that the economic actors do know these parameters perfectly. This approach focuses on the uncertainty of the "modeler" and observer, rather than uncertainty faced by the "decision maker". It is unclear if the analysis of parametric uncertainty should rely on certainty equivalent formulations, which in many contexts are not reliable analyses of how risk-averse agents respond to uncertainties in a dynamic world.

A different approach is taken in a much smaller part of the literature focusing on uncertainty, i.e. the realization of the world is uncertain due to some random events. See e.g. Kelly and Kolstad (1999) and Crost and Traeger (2011). However, the IAM community has so far not produced a stochastic IAM flexible enough to represent uncertainty in a quantitatively realistic manner. Most existing stochastic dynamic optimization approaches to model uncertainty in IAM suffer from one or more of the following features: 1) long duration of time periods (e.g., 5 or 10 years), 2) short time horizon or small number of periods (e.g., 2 periods), 3) reduced dimension of a full scale IAM (e.g., a two-dimensional version of the seven-dimensional model), and finally 4) expandability of the model (e.g., the flexibility of the model and computational method dealing with a deepened complexity of the state space).

Several studies build on DICE2007 (Nordhaus, 2008) and attempt to study optimal climate policies under uncertainty. Examples include Kelly and Kolstad (1999) and Bahn et al. (2008). However, just as in the DICE2007 model, these studies assume 10-year time units. We argue that the representation of time should be compatible with the frequencies of both the natural and social processes related to climate change. Any IAM with 10-year time periods represent neither social nor physical processes because nontrivial dynamics and feedbacks may occur in either system during a single decade. Dynamic stochastic general equilibrium (DSGE) models in economics use relatively short time periods, always at most a year. Decadal time periods as in DICE2007 are too long for serious, quantitative analysis of policy questions. For example, if one wants to know how carbon prices should react to business cycle shocks, the time period needs to be at most a year. No one would accept a policy that takes ten years to respond to current shocks to economic conditions. Cai, Judd and Lontzek (2012c) shows that annual time periods produce a significantly different carbon price with the numbers given by DICE2007. Cai, Judd and Lontzek (2012a) develops DICE-CJL, a continuous-time extension of DICE2007, and then demonstrate that many substantive results depend critically on the time step, strongly supporting our contention that short time periods are necessary for a quantitatively reliable analysis.

Other IAM analyses that examine uncertainty assume a very small number of decision periods, often with very large time periods. For example, Webster et al. (2012) models a 7-period horizon version of DICE with 50-year time units, while the model of Fisher and Narain (2003) has only 2 periods. In general, while a model with a small number of decision periods might deliver some important qualitative insights of the model, it is highly unlikely that these models can capture the full dynamic evolution of the climate and economic systems.

We argue that any reliable IAM should incorporate the ability to study the optimal decision making over a very long time horizon with several hundred decision periods. A stochastic formulation of such a model necessitates a solution method using dynamic programming. As of today there are only a handful of models following that approach. Lontzek and Narita (2011) study a continuous-time, infinite horizon model integrating climate and the economy. A recent study by Lemoine and Traeger (2012) applies discrete-time dynamic programming methods. Both models build on DICE2007 but examine a lower-dimensional climate system for reasons of computational tractability. The DICE's six-dimensional state transition system is one in which all elements interact and there is no lower dimensional system that is equivalent. Therefore, we argue that, except for a measure zero set of initial conditions and impulses, a lower-dimensional representation of a full dynamic system cannot capture the full spectrum of the DICE results. Kelly and Kolstad (1999) also implements risk. However, the methodological approach in Kelly and Kolstad (1999) makes it very difficult to extend the analysis to a higher-dimensional space and time frequency.

It is highly desirable to produce a DSGE IAM, incorporating stochastic evolution of the climate and economic system. However, it is often argued that a complete stochastic analysis of IAMs is not possible due to computational complexity. In fact, in a recent report (EPA, 2010) the Environmental Protection Agency assesses that currently the Integration of DSGE models with long-run inter temporal models, e.g., IGEM (Goettle et al. 2009), is beyond the scientific frontier.

A truly stochastic model would also enable a sound analysis of abrupt and irreversible climate change which so far has only been studied in very simple models. Full-size versions of IAMs have only dealt with this topic by using probabilistic assessment studies. For example, Nordhaus (2008) uses a certainty-equivalent approach to tipping points and argues that the additional carbon tax resulting from tipping points exhibits a ramp structure, i.e., it is initially mild but rises substantially over time with global warming. A recent study by Lontzek, Cai and Judd (2012) uses the same model as Nordhaus (2008) but uses a stochastic formulation of abrupt and irreversible climate change and obtains a completely different result. The possibility of a low-probability and low-impact tipping event results in a flat profile for the additional carbon tax. This example demonstrates how implications for climate policy depend on the appropriate model formulation of the underlying research question. It is reasonable to assume that a fully stochastic version of any IAM will unveil new and interesting results.

Novel findings for the SCC can be expected from using more flexible alternatives to the standard preference specifications. In particular, since riskiness is inherent in the nature of tipping points, the socio-economic effects of stochastic abrupt and irreversible climate change will be affected by preferences about risks. Hence, it is natural to combine Epstein-Zin preferences (Kreps and Porteus 1978, Epstein and Zin 1989) with tipping point risks. It seems likely, that including risk-sensitive preferences will surely imply greater willingness to pay to avoid adverse climate change.

For this study we use DSICE (Cai, Judd and Lontzek, 2012b), a DSGE full-dimensional extension of DICE2007. DSICE and DICE2007 are therefore comparable which facilitates the comparison of our carbon tax numbers with the ones obtained, e.g., by the U.S. Government Interagency Working Group on Social Cost of Carbon (IWG, 2010). We include Epstein-Zin preferences as well as a system of tipping points. We are thus endowed with a model structure that has a more realistic description of risk preferences, the stochastic processes of damages and business cycle fluctuations. We model abrupt and irreversible climate change as a low-probability and low-impact event. Thus, we abstain from dramatic catastrophe assumptions and follow rather a conservative approach. We also assume business cycle shocks that are moderate and bounded. We solve DSICE with dynamic programming using advanced computational methods. The solution to DSICE is reliable and quickly obtained.

As in Lontzek, Cai and Judd (2012), we find that the threat of a tipping point induces immediately stringent carbon pricing, even for low-probability and low-impact tipping events. We find that including Epstein-Zin preferences into DSICE significantly increases the carbon tax. For example, larger values of the intertemporal elasticity of substitution and the degree of risk aversion lead to higher carbon taxes. Furthermore, in addition to abrupt and irreversible climate change, we also study cases with significant uncertainty about the post-tipping damage impact. We find that the uncertainty about the damage is also a critical factor leading to a sharp increase in the carbon tax. Furthermore, for low degrees of risk aversion (i.e. smaller than 2), the carbon tax is not affected much by the volatility of the uncertain damage. In contrast, high degrees of risk aversion significantly amplify the effect of damage uncertainty on the carbon tax. We also investigate a disaster scenario. We find that if there is a very unlikely tipping point event (about 0.1% probability of tipping at 2100), and it is has major permanent but uncertain impact on productivity (mean is 20% and volatility is 10%), today's carbon tax of a highly risk-averse policy maker is $1170.

Besides the basic structure of the DSICE model, this study's analysis includes four major components: 1) the modeling of abruptness and irreversibility of climate change, 2) the modeling of non-separable preferences, 3) the modeling of business cycle uncertainty and 4) the numerical algorithm and the applied computational methods. We proceed as follows: Sections 2 and 3 discuss our modeling of the stochastic climate, stochastic economy and the preferences. Section 4 presents the DP framework of the DSICE model with Epstein-Zin preferences, and then introduces the multidimensional numerical dynamic programming method we use to solve it. Section 5 presents our results and Section 6 concludes.

2 Abrupt and Irreversible Climate Change

In order to implement the risk of tipping points into an IAM, we need to model a hazard rate for each possible tipping point event. Thus, for any given combination of the state space in any year, the hazard rate gives us the probability of a particular tipping point occurring in that year. Consequently, the tipping point event becomes stochastic.

So far, most studies model abrupt climate change as a deterministic process. For example, Mastrandrea and Schneider (2001) couple an older version of the DICE model with a simple climate model which simulates the functioning of the THC. Information obtained from the additional climate model is used to enhance the exponential component in the DICE damage function. This specification gives rise to a steeper damage function (in terms of global warming) but is inconsistent with the nature of the externality of an irreversible catastrophe. Nordhaus (2010) also models sea level rise deterministically by increasing the curvature of the damage function.

Another branch of the literature on catastrophic climate change in IAMs (e.g., Yohe et al. 2004) models tipping point events as happening when a threshold temperature is passed (or some other condition is met). In our view this approach (or at least the simplest version of it) is not appealing because it implies that if we have been at that threshold temperature level, then we can immediately infer (learn) that we are safe as long as we stay at or below that level. Modeling abrupt climate change with a known threshold location is a special case of our hazard function approach. In that special case (e.g., Keller et al., 2004) the hazard rate equals zero up to the threshold temperature level and is equal to one beyond that level. In contrast, knowing that there is a critical point but not knowing where it is would not imply a hazard rate. The latter formulation implies that once we have reached a temperature level with no tipping event, we conclude that there cannot be a tipping point at that temperature level. This is not in the statistical nature of a hazard rate.

Instead we assume that even if a tipping point has not yet occurred, it may still occur later even if there is no further warming, or even if there is cooling. This surely will affect the optimal climate policy. For example, the optimal mitigation policy in DICE2007 increases rapidly over time, with small efforts today but much greater effort at the end of the current century. This “ramp” structure of mitigation policy implied by DICE2007-style models arises because of the nature of the externality, i.e., a high atmospheric temperature today reduces output today, but there is no direct impact on future damages, which depend only on contemporaneous temperatures. However, by including the possibility of an abrupt and irreversible climate change, one actually changes the nature of the climate externality. Figure 1 highlights the two aspects of the global warming externality.

Figure 1: The nature of externality from abrupt and irreversible climate change

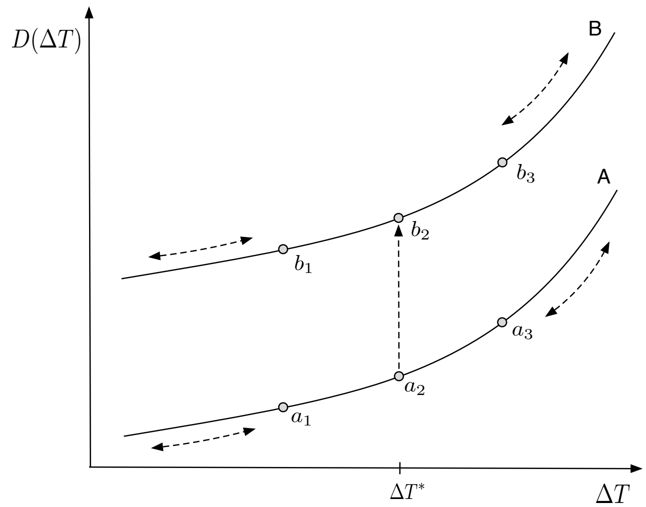

Curve A represents the pre-tipping damage factor as a function of temperature change (ΔT). It is monotonically increasing and convex. In a pre-tipping regime any movement along A (i.e., $a_1 \leftrightarrow a_2 \leftrightarrow a_3$) can occur, depending on the change in global average temperature. As a consequence, the resulting damage rate $D(\Delta T)$ will exhibit a smooth pattern. The same logic applies in a post-tipping regime, which is denoted by B. Any movement along B (i.e., $b_1 \leftrightarrow b_2 \leftrightarrow b_3$) can occur, depending on the change in global average temperature after a tipping point has occurred. Now, assume that a tipping point has occurred at a_2, which corresponds to a global warming level of ΔT^* [1]. In that case, the damage rate resulting from the change in the climate system will be represented by B (i.e., $a_2 \rightarrow b_2$). However, the irreversible nature of abrupt climate change prevents a movement back to A in case of global cooling and temperature levels below the realized tipping point level. Instead the damage factor will remain to move along B. This stochastic aspect of a tipping point can only be captured by specifying a hazard rate and modeling abrupt climate change as a jump process where the Markovian hazard rate only depends on contemporaneous conditions. It cannot be captured by changing the shape (e.g., increasing the exponents) of the damage function of global warming.

Unfortunately, as Kriegler et al. (2009) points out, the lack of data and limited understanding of the underlying processes make it difficult to assess the likelihood of changes in the earth system due to global warming. As of today, most probability assessments of abrupt climate change come from expert elicitation surveys (e.g., Kriegler et al. 2009 and Zickfeld et al. 2007). In these elicitations, the experts are asked about their beliefs about probabilities of a tipping point occurring at, or before

[1] Note that our specification allows for a tipping point to occur even in a phase of global cooling. This feature results from the stochastic nature of a hazard rate formulation of abrupt climate change.

some time (typically the year 2100 or 2200) based on an emission scenario. These experts' opinions reveal their probabilistic beliefs for a tipping point occurring within a time frame under alternative assumptions about temperature changes within that specific time frame.

Zickfeld et al. (2007) presents experts' subjective cumulative probabilities that a collapse of THC will occur or be irreversibly triggered at or before 2100 under alternative assumptions about temperature at 2100. The survey data obtained in Zickfeld et al. (2007) on THC collapse reveal a huge range of assigned probabilities. For example, some of these experts believe that even with an additional warming of $6°C$ it is impossible for the THC to reach a tipping point by 2100. On the contrary, one third of the experts believe that the probability is $> 60\%$ and two experts even suggest 90%.[2] For less global warming by 2100 these elicited probabilities are accordingly lower. The average expert assigns about 18% chance of tipping at $4°C$ by 2100 and 4% at $2°C$. A similar large range of experts' opinions is obtained by Kriegler et al. (2009) which presents the tipping probabilities for five tipping elements. For the medium temperature corridor (i.e. a warming of $2-4°C$ by 2200), this study finds a range of roughly 50% for the THC, 80% for GIS, WAIS, and AMAZ, and about 40% for ENSO. Overall, the numbers from both studies mirror a remarkably huge range of expert opinions about the tipping point probabilities. This incongruity in probability assessment among the experts makes the analysis of a tipping point even more interesting from a uncertainty point of view.

In order to derive the hazard rates for our abrupt climate change examples, we use the intrinsic information provided in the expert elicitation studies by Zickfeld et al. (2007) and Kriegler et al. (2009). We follow the approach in Lontzek, Cai and Judd (2012) which is the first study to incorporate the results of expert elicitation study into a stochastic IAM. The approach in Lontzek, Cai and Judd (2012) is to use the experts' opinions on the cumulative probability of a tipping point occurring by 2100 or 2200, and infer hazard rates for tipping elements from these opinions and their assumptions about the temperature path. The resulting hazard rate becomes a function of the contemporaneous temperature level (measured as the deviation from preindustrial temperature).

3 A Stochastic IAM with Epstein-Zin Preferences

The DSICE model (Cai, Judd and Lontzek, 2012b) is a dynamic stochastic general equilibrium model integrating climate and the economy. DSICE is basically a DGSE-extension of the DICE-CJL model (Cai, Judd and Lontzek, 2012a), which itself is a numerically stable version of DICE2007 with a flexible time-period length. This version of DSICE incorporates both, uncertainty about the future state of climate and the economy. A complete description of the DICE-CJL model with all equations as well as parameters can be found in Cai, Judd and Lontzek (2012a). Furthermore, the

[2] The probability numbers have been eyeballed from Zickfeld et al. (2007).

DSICE model with separable utility is described in Cai, Judd and Lontzek (2012b). In this section we briefly describe the DSICE model focusing mainly on Epstein-Zin preferences, the stochastic climate and the stochastic economy.

3.1 Epstein-Zin Preferences

The standard separable utility function in the finite-horizon DICE2007-class of models is

$$u(c_t, l_t) = \frac{(c_t/l_t)^{1-\psi}}{1-\psi} l_t,$$

where c_t denotes the consumption level and l_t is total labor supply. It is assumed that a social planner maximizes the present-discounted utility stream up to a terminal time T. In order to formulate the utility function with Epstein-Zin preferences into DSICE we account for the 7-dimensional continuous state space of the model. The later consist of k, the capital stock, $\mathbf{M}$, the three-dimensional carbon system, $\mathbf{T}$, the two-dimensional temperature vector and ζ the economic shock. Furthermore, J is the discrete shock to the climate. Given that $U_t = \sum_{t=0}^{T} \beta^t u_t(c_t, l_t)$, a recursive formulation of the social planner's maximand is

$$\begin{aligned} U_t\left(k, \mathbf{M}, \mathbf{T}, \zeta, J\right) &= \max_{c,\mu} \left\{ (1-\beta) \frac{(c_t/l_t)^{1-\psi}}{1-\psi} l_t + \right. \\ &\left. \beta \left[\mathbb{E}\left\{ \left(U_{t+1}\left(k^+, \mathbf{M}^+, \mathbf{T}^+, \zeta^+, J^+\right)\right)^{1-\gamma} \right\} \right]^{\frac{1-\psi}{1-\gamma}} \right\}^{\frac{1}{1-\psi}} \end{aligned}$$

where ψ is the inverse of the intertemporal elasticity of substitution, μ_t is the mitigation rate of emissions, $\mathbb{E}\{\cdot\}$ is the expectation operator, β is the discount factor and γ is the degree of risk aversion.

Epstein-Zin preferences are flexible specifications of decision makers' preferences regarding uncertainty. According to Epstein-Zin specifications about preferences, there is value in resolving some of the uncertainty and we are interested in computing the real willingness to pay to reduce risks by exploring ranges of parameters, consistent with existing risk premium. We know from the literature on equity premium that the willingness to pay to reduce risk is far higher than implied by standard utility functions, and therefore we expect an increase in the optimal carbon price due to tipping point risk. As of today, Epstein-Zin preferences have not been analyzed in a full-dimensional DSGE version of a major IAM, such as DICE2007. On top of this, we account for an additional stochastic dimension by modeling abrupt and irreversible climate change with a system of tipping elements.

3.2 The Stochastic Climate

Let $\mathbf{M}_t = (M_t^{\mathrm{AT}}, M_t^{\mathrm{UP}}, M_t^{\mathrm{LO}})^\top$ be a three-dimensional vector describing the masses of carbon concentrations in the atmosphere, and upper and lower levels of the ocean.

These concentrations evolve over time according to:

$$\mathbf{M}_{t+1} = \mathbf{\Phi}^{\mathrm{M}}\mathbf{M}_t + (\mathcal{E}_t, 0, 0)^\top,$$

where

$$\mathbf{\Phi}^{\mathrm{M}} = \begin{bmatrix} 1-\phi_{12} & \phi_{12}\varphi_1 & 0 \\ \phi_{12} & 1-\phi_{12}\varphi_1-\phi_{23} & \phi_{23}\varphi_2 \\ 0 & \phi_{23} & 1-\phi_{23}\varphi_2 \end{bmatrix},$$

with $\varphi_1 = M_*^{\mathrm{AT}}/M_*^{\mathrm{UP}}$ and $\varphi_2 = M_*^{\mathrm{UP}}/M_*^{\mathrm{LO}}$, where M_*^{AT}, M_*^{UP} and M_*^{LO} are the preindustrial equilibrium states of the carbon cycle system. The anthropogenic sources of carbon are represented by the $\mathcal{E}_t$, which will be specified in the next subsection.

The DICE climate system also includes temperatures in the atmosphere and ocean, which are represented by the vector $\mathbf{T}_t = (T_t^{\mathrm{AT}}, T_t^{\mathrm{LO}})^\top$. The temperatures dynamically evolve according to:

$$\mathbf{T}_{t+1} = \mathbf{\Phi}^{\mathrm{T}}\mathbf{T}_t + \left(\xi_1\mathcal{F}_t\left(M_t^{\mathrm{AT}}\right), 0\right)^\top,$$

where the heat diffusion process between ocean and air is represented by the matrix

$$\mathbf{\Phi}^{\mathrm{T}} = \begin{bmatrix} 1-\xi_1\eta/\xi_2-\xi_1\xi_3 & \xi_1\xi_3 \\ \xi_4 & 1-\xi_4 \end{bmatrix},$$

where ξ_2 is the climate sensitivity parameter (we choose $\xi_2 = 3$ in our examples of this paper). Atmospheric temperature is affected by external forcing, F_t^{EX}, and by the interaction between radiation and atmospheric CO2, implying that total radiative forcing at t is

$$\mathcal{F}_t\left(M^{\mathrm{AT}}\right) = \eta\log_2\left(M^{\mathrm{AT}}/M_0^{\mathrm{AT}}\right) + F_t^{\mathrm{EX}}. \tag{1}$$

The impact of global warming on the economy is reflected by a convex damage function of temperature in the atmosphere. This is a standard feature of the DICE model family. As discussed in the introduction, we modify the standard damage function by explicitly modeling the possibility of a climate shock (i.e. tipping point) to account for the threat of abrupt and irreversible climate change[3]. Each climate shock occurs at a random time. Thus, the stochastic damage function in DSICE is given by

$$\Omega\left(T_t^{\mathrm{AT}}, J_t\right) = \frac{J_t}{1+\pi_1 T_t^{\mathrm{AT}} + \pi_2 (T_t^{\mathrm{AT}})^2},$$

where the denominator represents the standard damage function and the numerator J_t, a discrete Markov chain with 2 possible values. $J_t = 1$ in the pre-tipping regime

[3] At this stage we only describe one single tipping point. However, we also analyze multiple tipping points as well as multi-stage tipping points. These examples are discussed later in this study

and $0 < J_t < 1$ in the post-tipping regime. The Markov chain probability transition matrix at time t is

$$\begin{bmatrix} 1-p_t & p_t \\ 0 & 1 \end{bmatrix}$$

where its (i,j) element is the transition probability from state i to j for J_t, and p_t is the conditional probability that a tipping event occurs at time t, and depends positively on surface temperature at time t. State 2 is an absorbing state, representing the irreversible nature of the tipping point.

3.3 The Stochastic Economy

Capital k_t transits to the next period in a standard fashion:

$$k_{t+1} = (1-\delta)k_t + \mathcal{Y}_t(k_t, T_t^{\mathrm{AT}}, \mu_t, \zeta_t, J_t) - c_t,$$

where $\mathcal{Y}_t$ denotes the stochastic production function. The latter accounts for the costs of mitigation as a fraction of output. Furthermore, it includes the damage resulting from global warming as well as both, an economic shock and a climate shock:

$$\mathcal{Y}_t(k_t, T_t^{\mathrm{AT}}, \mu_t, \zeta_t, J_t) = (1-\theta_{1,t}\mu_t^{\theta_2})\zeta_t A_t k_t^{\alpha} l_t^{1-\alpha} \Omega\left(T_t^{\mathrm{AT}}, J_t\right),$$

where ζ_t is a discrete-time bounded mean-reverting continuous productivity shock representing economic fluctuations (see Cai, Judd and Lontzek 2012b), and its transition function from stage t to $t+1$ is $\zeta_{t+1} = g^{\zeta}(\zeta_t, \omega_t^{\zeta})$ where ω_t^{ζ} is an i.i.d. random process. We assume that the economic shock and the climate shock are independent.

Given the stochastic production function, annual total carbon emissions are stochastic and given by

$$\mathcal{E}_t\left(k_t, \mu_t, \zeta_t\right) = \sigma_t(1-\mu_t)\zeta_t A_t k_t^{\alpha} l_t^{1-\alpha} + E_t^{\mathrm{Land}}, \tag{2}$$

where σ_t denotes the carbon intensity of output, μ_t denotes the fraction of mitigated emission and E_t^{Land} is an exogenous rate of emissions from biological processes.

4 The Dynamic Programing Problem

We solve the stochastic Integrated Assessment Model with Epstein-Zin preferences numerically. This section shows the dynamic programing formulation of the model. Furthermore, we present the numerical algorithm and specifically focus on some essential aspects of our computational method.

4.1 DP formulation

Using the recursive utility function and letting

$$V_t\left(k, \mathbf{M}, \mathbf{T}, \zeta, J\right) = \frac{\left[U_t\left(k, \mathbf{M}, \mathbf{T}, \zeta, J\right)\right]^{1-\psi}}{\left(1-\psi\right)\left(1-\beta\right)}$$

denote the value function we can write the dynamic programing model for DSICE with Epstein-Zin preferences. Assuming

$$u(c, l_t) = \frac{\left(c_t/l_t\right)^{1-\psi}}{1-\psi} l_t,$$

the DP problem becomes

$$\begin{aligned}
V_t(k, \mathbf{M}, \mathbf{T}, \zeta, J) = \max_{c,\mu} \quad & u(c_t, l_t) + \frac{\beta}{1-\psi} \times \\
& \left[\mathbb{E}\left\{\left(\left(1-\psi\right) V_{t+1}\left(k^+, \mathbf{M}^+, \mathbf{T}^+, \zeta^+, J^+\right)\right)^{\frac{1-\gamma}{1-\psi}}\right\}\right]^{\frac{1-\psi}{1-\gamma}}, \\
\text{s.t.} \quad & k^+ = (1-\delta)k_t + \mathcal{Y}_t(k, T^{\mathrm{AT}}, \mu, \zeta, J) - c_t, \\
& \mathbf{M}^+ = \mathbf{\Phi}^{\mathrm{M}}\mathbf{M} + \left(\mathcal{E}_t\left(k, \mu, \zeta\right), 0, 0\right)^\top, \\
& \mathbf{T}^+ = \mathbf{\Phi}^{\mathrm{T}}\mathbf{T} + \left(\xi_1 \mathcal{F}_t\left(M^{\mathrm{AT}}\right), 0\right)^\top, \\
& \zeta^+ = g^\zeta(\zeta, \omega^\zeta), \\
& J^+ = g^J(J, \mathbf{T}, \omega^J),
\end{aligned}$$

for $t = 0, 1, \ldots, 599$, and the terminal value function V_{600} is the terminal value function given in Cai, Judd and Lontzek (2012b). In the model, consumption c and emission control rate μ are two control variables, $(k, \mathbf{M}, \mathbf{T}, \zeta, J)$ is 8-dimensional state vector at year t (where $\mathbf{M} = (M^{\mathrm{AT}}, M^{\mathrm{UP}}, M^{\mathrm{LO}})^\top$ is the three-layer CO_2 concentration and $\mathbf{T} = (T^{\mathrm{AT}}, T^{\mathrm{LO}})^\top$ is the two-layer global mean temperature), and $(k^+, \mathbf{M}^+, \mathbf{T}^+, \zeta^+, J^+)$ is its next-year state vector.

4.2 The Dynamic Programming Algorithm

In dynamic programming problems, when the value function is continuous, it has to be approximated. In this study, we use a finitely parameterized collection of functions to approximate a value function, $V(x, \theta) \approx \hat{V}(x, \theta; \mathbf{b})$, where x is the continuous state vector (in DSICE, it is the 7-dimensional vector $(k, \mathbf{M}, \mathbf{T}, \zeta)$, θ is the discrete state vector (in DSICE, it is J), and $\mathbf{b}$ is a vector of parameters. The functional form $\hat{V}$ may be a linear combination of polynomials, or it may represent a rational function or neural network representation, or it may be some other parameterization especially designed for the problem. After the functional form is fixed, we focus on finding the vector of parameters, $\mathbf{b}$, such that $\hat{V}(x, \theta; \mathbf{b})$ approximately satisfies the Bellman equation. Numerical DP with value function iteration can solve the Bellman equation approximately (Judd, 1998). A general DP model is based on the Bellman equation:

$$
\begin{aligned}
V_t(x,\theta) = \max_{a\in\mathcal{D}(x,\theta,t)} \quad & u_t(x,a) + \beta\mathbb{E}\left\{V_{t+1}(x^+,\theta^+)\right\}, \\
\text{s.t.} \quad & x^+ = f(x,\theta,a), \\
& \theta^+ = g(x,\theta,\omega),
\end{aligned}
$$

where $V_t(x,\theta)$ is the value function at time $t \le T$ (the terminal value function $V_T(x,\theta)$ is given), (x^+,θ^+) is the next-stage state, $\mathcal{D}(x,\theta,t)$ is a feasible set of a, ω is a random variable, β is a discount factor and $u_t(x,a)$ is the utility function at time t. The following is the algorithm of parametric DP with value function iteration for finite horizon problems. Detailed discussion of numerical DP can be found in Cai (2009), Judd (1998) and Rust (2008).

Algorithm 1. *Numerical Dynamic Programming with Value Function Iteration for Finite Horizon Problems*

Initialization. *Choose the approximation nodes,* $X_t = \{x_{it} : 1 \le i \le m_t\}$ *for every* $t < T$*, and choose a functional form for* $\hat{V}(x,\theta;\mathbf{b})$*, where* $\theta \in \Theta$*. Let* $\hat{V}(x,\theta;\mathbf{b}_T) \equiv V_T(x,\theta)$*. Then for* $t = T-1, T-2, \ldots, 0$*, iterate through steps 1 and 2.*

Step *1. Maximization step. Compute*

$$
\begin{aligned}
v_{i,j} = \max_{a_{i,j}\in\mathcal{D}(x_i,\theta_j,t)} \quad & u_t(x_i,a_{ij}) + \beta\mathbb{E}\left\{\hat{V}(x^+_{i,j},\theta^+_j;\mathbf{b}_{t+1})\right\} \\
\text{s.t.} \quad & x^+_{i,j} = f(x_i,\theta_j,a_{i,j}), \\
& \theta^+_j = g(x_i,\theta_j,\omega),
\end{aligned}
$$

for each $\theta_j \in \Theta$*,* $x_i \in X_t$*,* $1 \le i \le m_t$*.*

Step *2. Fitting step. Using an appropriate approximation method, compute the* $\mathbf{b}_t$ *such that* $\hat{V}(x,\theta_j;\mathbf{b}_t)$ *approximates* $(x_i, v_{i,j})$ *data for each* $\theta_j \in \Theta$*.*

We implement our numerical dynamic programming algorithm to solve the DSICE model with Epstein-Zin preferences. The code is written in Fortran and uses the methods presented in Judd (1998), Cai (2009), and Cai and Judd (2010, 2012a, 2012b, 2012c), and we use NPSOL (Gill, P., et al., 1994) as the optimization solver in the maximization step. In particular, for each discrete state value, we choose the degree four complete Chebyshev polynomials to approximate the value function. The multidimensional tensor grid with five Chebyshev nodes on each continuous dimension in the continuous state ranges gives us our set of approximation nodes. The continuous state ranges are constructed from the solution of the DICE-CJL model (Cai, Judd and Lontzek, 2012a), a continuous-time reformation of DICE2007

(Nordhaus, 2008). The results of the maximization step at each of the approximation nodes are used to compute the Chebyshev coefficients via a regression procedure. We assume a one-year time period and iterate backwards for 600 years, beginning with a terminal value function. See Cai, Judd and Lontzek (2012b) for a description of how we construct the terminal value function and the continuous state ranges and how we verify the accuracy of solutions given by our algorithm.

It is important to use one year instead of ten years as the time unit. Cai, Judd and Lontzek (2012a, 2012c) shows that annual time periods produce a significantly different carbon price with the numbers given by DICE2007 using ten-year time periods. For the models with uncertain state variables, Cai, Judd and Lontzek (2012b) also shows that annual time periods have a significantly higher carbon price (about 24% higher in the first period) than ten-year time periods.

After computing the approximate value function at each time t for $t = 0, 1, ..., 599$, we simulate the optimal path corresponding to a sequence of shocks, both economic and climate. That is, given the current state along a path, we compute the optimal decisions and then use the realized shock to compute next period's state. We start this process with the given initial continuous state and $(\zeta_0, J_0) = (1, 1)$, and run it until the terminal time. The next section presents the statistical results of 1000 simulated paths of the optimal solution.

5 Results and Discussion

Here, we analyze how the optimal carbon tax is affected by different preference parameters combined with various tipping point events. The parameters of the deterministic part of DSICE are the same as in the original DICE2007 model and described in Cai, Judd and Lontzek (2012) in more detail. Thus, we stay as close as possible to the original DICE2007 model. Regarding the preference parameters, the degree of risk aversion γ and the reciprocal of the intertemporal elasticity of substitution ψ, we cover a broader range for both in accordance with empirical studies. Basal and Yaron (2004), e.g., finds $\gamma = 10$ and $\psi = 2/3$. Vissing-Jørgensen and Attanasio (2003) finds γ between 5 and 10 and $\psi < 1$, while Vissing-Jørgensen (2002) and Campbell and Cochrane (1999) also find cases of $\psi > 1$. Furthermore, Nordhaus (2008) chooses $\psi = 2$ (which implies that $\gamma = 2$). We therefore choose a range for γ between 0.5 and 20 and for ψ between 0.5 and 2.

5.1 DSICE with One Tipping Event

We first examine the case of one tipping event. For illustrative purposes, we first do not account for business cycle shocks, which we include in a next step. Recall, that the probability transition matrix of J_t from year t to year $t + 1$ is

$$\begin{bmatrix} 1 - p_t & p_t \\ 0 & 1 \end{bmatrix},$$

where its (i, j) element is the transition probability from state i to j. The probability p_t depends positively on the surface temperature at time t.

$$p_t = 1 - \exp\left\{-\nu \max\left\{0,\ (T_t^{\mathrm{AT}} - 1)\right\}\right\}, \tag{3}$$

where ν is called the hazard rate parameter. As discussed in Section 2, we specify the hazard rates for the tipping events in this study based on expert elicitation studies by Zickfeld et al. (2007) and Kriegler et al. (2009). The hazard rates in our numerical examples are motivated by the range of hazard rates that experts put on the various tipping events. Our objective is not to model any specifically tipping element explicitly but rather to show how that range of numbers is related to the SCC.

Figure 2: Carbon Tax in the benchmark scenario

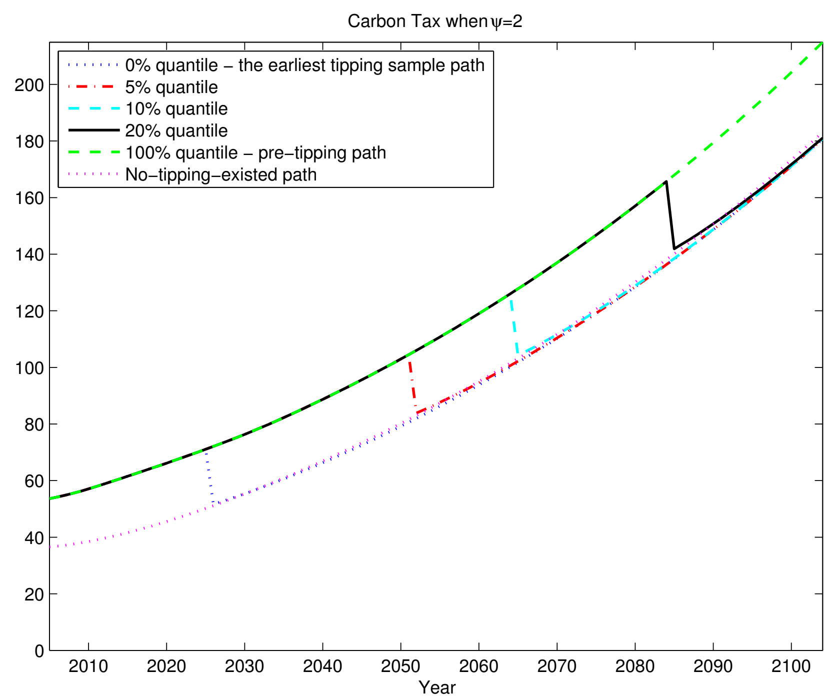

Figure 2 shows the statistical results of the simulation runs for the optimal carbon tax policy for our baseline scenario: $\psi = 2$, $\gamma = 10$, the hazard rate parameter $\nu = 0.00574$, and the damage level is 2.5% (i.e., $J_t = 0.975$ when the tipping event happens). We depict the pre-tipping path, the earliest tipping sample path as well as some relevant quantiles. Furthermore, we depict the basic deterministic carbon tax scenario in which a tipping point does not exist. Thus, we can compare the cases in which the economy is threatened by a tipping point with the case in which no tipping point (and the associated risk and impact) exists. The upper envelope in Figure 2 represents at each time t, the carbon tax if there has not yet been a tipping event. We call this the pre-tipping carbon tax. In contrast, the lower envelope in general represents the carbon tax in the post-tipping regime. Figure 2 also displays the timing of some sample tipping events. For example, the first vertical drop (which is at about 2027 in Figure 2) is the first tipping out of our 1000 simulations. By the middle of this century about 5% of the simulated paths have generated a tipping point and by the end of the 21st century more than 20% of the paths have exhibited a tipping point. Furthermore, note that in the initial period (2005), the optimal carbon tax is \$54, while it is \$37 in the case when the tipping point does not exist. Thus, in face of a low probability - low impact event the immediate additional preventive carbon tax is \$17.

Figure 2 also resembles the two distinct aspects of the abrupt climate change externality which have been discussed by Lontzek, Cai and Judd (2012). On the one hand, the DICE-like "ramp" structure of the carbon tax implies a rather low carbon tax for low global warming but then it becomes more stringent over time when global warming becomes more severe. This is shown by the no-tipping path in Figure 2. On the other hand, the threat of an irreversible and abrupt climate change results in a nearly constant additional carbon tax (i.e., the "ramp" is not the dominant feature.) to delay the tipping point from occurring. This result is derived from the drop in the carbon tax immediately after the tipping point event. Figure 2 implies a carbon tax markup of about 25%. As in Lontzek, Cai and Judd (2012), we conclude that the optimal carbon tax response to the threat of abrupt and irreversible climate change depends on the dynamic pattern of the adverse impacts. If these impacts are permanent, as it is the case here, the optimal policy is one with substantial carbon taxation immediately.

The results in Figure 2 are based on the benchmark Epstein-Zin parameter, $\psi = 2$ and $\gamma = 10$. We have argued previously that riskiness is inherent in the nature of tipping points. In the following, we therefore analyze how the optimal carbon tax is affected by stochastic abrupt and irreversible climate change under different preferences about risk and inter temporal substitution. Furthermore, we conduct some other sensitivity analyses.

5.1.1 Effect of Intertemporal Elasticity of Substitution

First, we consider the effect of a higher intertemporal elasticity of substitution on the carbon tax and analyze the cases with $\psi = 1.2$ and $\psi = 0.5$. The general finding from Figure 3 is that the carbon tax will increase as the intertemporal elasticity increases (i.e., ψ decreases). This can be best seen in the initial period, where the carbon tax almost doubles when moving from $\psi = 2$ to $\psi = 1.2$.

Figure 3: Carbon Tax for different intertemporal Elasticity

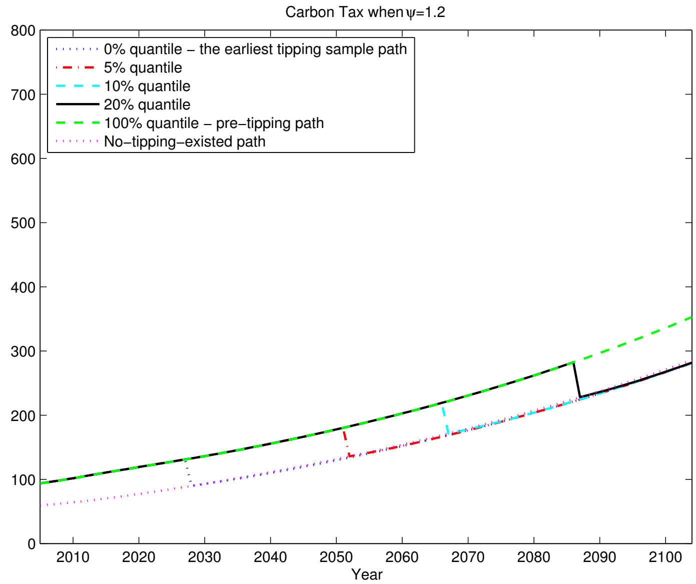

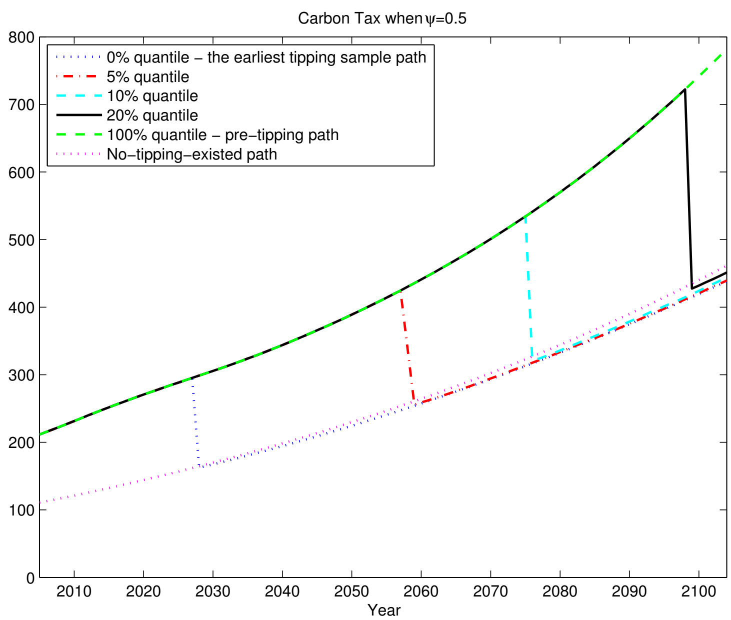

A further doubling of the tax in the initial period occurs when going from $\psi = 1.2$ to $\psi = 0.5$. Moreover, the difference between the pre-tipping path (100% quantile line) and the post-tipping path (0% quantile line) is increasing as the intertemporal elasticity increases. This additional preventive carbon tax around 2050 rises from about \$20 ($\psi = 2$) , to \$40 ($\psi = 1.2$) and to about \$160 ($\psi = 0.5$). The strong effect of the IES on the carbon tax occurs because a low intertemporal elasticity of substitution, by definition, enhances the consumption-smoothing choices of the decision maker. In DSICE, these choices are savings and mitigation. On the one hand, more capital is built up to compensate for a potential loss in disposable output in case of a tipping point. On the other hand, mitigation efforts are very high and as a result, very few carbon emissions are released into the atmosphere.

5.1.2 Effect of Hazard Rates and Damage Level

While the previous example showed how the intertemporal elasticity of substitution affects the carbon tax, this section studies how the carbon tax is affected by different hazard rate parameters and damage levels. So far, we have considered a 2.5% impact of a tipping point on the economy. This is within the lower bound of what is assumed in the studies of catastrophic events. Only a few studies provide an estimate of the loss in output from a tipping point catastrophe, such as the THC collapse. Keller et al. (2004) estimate the loss in GDP from a tipping event in the range of 1% - 3%. Other studies, such as Mastrandrea and Schneider (2001) and Nordhaus (1994) use much higher estimates. In the DICE model catastrophic damages can amount up to 30% of GDP (Nordhaus, 2008) and in PAGE (Hope, 2006) up to 5%. To study the sensitivity of the carbon tax, we present the optimal carbon tax numbers for much higher impact levels on the economy. In particular, 5% and 10%. Lontzek, Cai and Judd (2012) have shown that the near constancy of the optimal anti-tipping effort is retained for different levels, and argued that this additional carbon tax does not depend on the magnitude of the damage but rather is inherent in the stochastic structure of the jump process shock. In this example, we retain the inverse of intertemporal elasticity of substitution at $\psi = 2$, and the risk-averse coefficient at $\gamma = 10$. Figure 4 shows the carbon tax at the initial year for different hazard rate parameters and damage levels. For example, the optimal carbon tax in face of a 10% damage and a hazard rate parameter of $\nu = 0.01$ is \$180, while it is only \$45 in face of a 2.5% damage and $\nu = 0.00228217$. Two general insights can be drawn from all four plots. First, the carbon tax increases more than proportionally with higher damage levels and second, the carbon tax also increases if the hazard rate parameter increases.

Figure 4: Effect of Hazard Rate Parameter and Damage Level

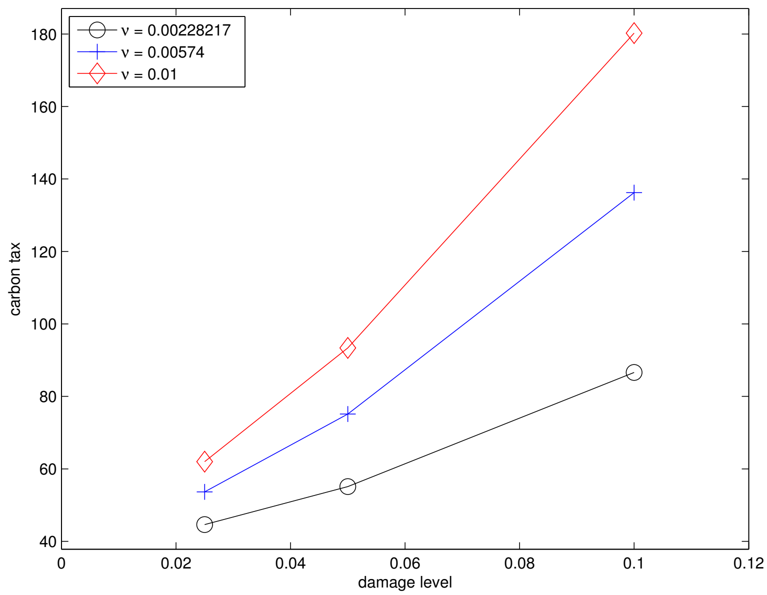

5.1.3 Putting a Price on a Disaster

So far we have considered tipping scenarios with a relatively low post-tipping impact. Recall that, in our benchmark scenario the post-tipping damage is 2.5% while our hazard rate parameter is $\nu = 0.00574$. We have observed the flattening of the additional carbon tax when the society accounts for a possible tipping point level. Our benchmark parameter specification resembles a "low probability - low impact" scenario for tipping points. Above, we also discussed cases with an impact level of up to 10% and saw that higher impact levels drastically raise the carbon tax. While an impact level of 10% is certainly not very low, it is still far below the levels of some disaster scenarios analyzed in the literature. As an example of a very high impact from a tipping point, we run the case for $J = 0.8$, i.e., a permanent 20% impact level. Figure 5 presents the optimal carbon tax in the face of a single tipping event resulting in a permanent damage of 20%. The carbon tax in the initial period is about $125, compared to our benchmark tax of about $54, which resembles the combined effect of much higher damages and a much lower hazard rate. This remarkably high increase in the carbon tax shows that the optimal climate policy towards catastrophic events implies a very high price for delaying (preventing) a low probability - low impact event. The pre-tipping carbon tax rises to about $400 by the end of this century and further exceeds $800 by 2160. Furthermore, note that the earliest sample tipping point occurs at around 2035, while by 2100 slightly more than 3% and by 2160 more than 10% of the sample paths exhibit a tipping point.

Figure 5: Disaster Case

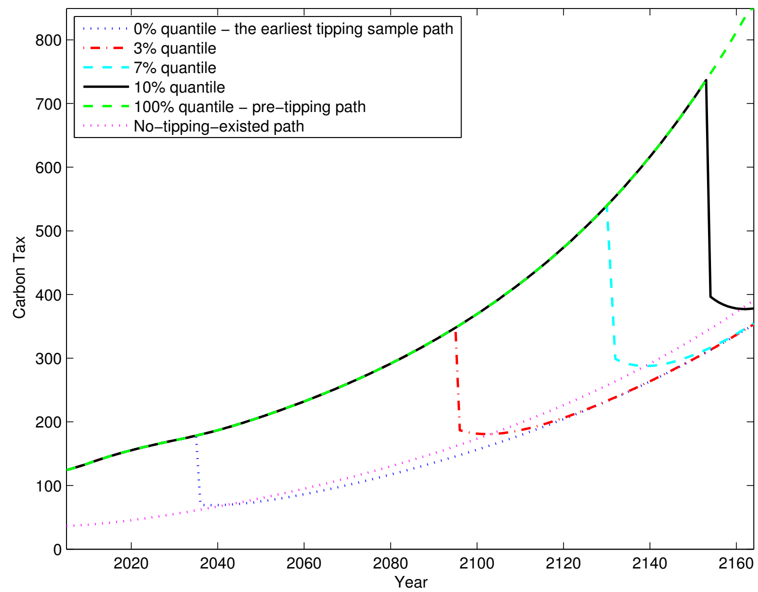

At the same time, we substantially reduce the hazard rate parameter to account for the most conservative lower assumptions about tipping point probabilities. In particular, we assume $\nu = 0.0008$, which results in a hazard rate of about 0.1% at 2100. This specification attributes a rather thin-tailed nature to our disaster case. Furthermore, we retain our benchmark preference parameters, $\psi = 2$, and $\gamma = 10$.

5.2 A Tipping Point with Uncertain Damage Level

We have argued before that specifying a number for the post-tipping damage is difficult. For one reason, it is a complex task to aggregate the multitude of sectoral and regional impacts. Another reason is that we don't know when a tipping point will occur in first place, and whether it occurs in 2050 or 2100, the impact will depend on many characteristics of the economy at that time. It is therefore reasonable to assume that the impact from a future tipping point is uncertain. This assumption, along with the uncertainty about the tipping point itself, creates an additional layer of uncertainty. In the following, we study cases in which both, the tipping point and the post-tipping magnitude of the damage are both uncertain. In the first example. we suppose that if the tipping event occurs, then

$$J_t = \begin{cases} 0.95 + \delta, & \text{with probability } 50\%, \\ 0.95 - \delta, & \text{with probability } 50\%, \end{cases} \tag{4}$$

where $0 \leq \delta < 0.05$ is called as the volatility of the uncertain damage level. Thus, we have three values for J_t, both of the post-tipping values are absorbing states, but there is now more risk that society faces given the tipping. We have both the risk of having a tipping, plus the unknown magnitude of the risk. Therefore, the tipping point has the probability transition matrix

$$\begin{bmatrix} 1-p_t & 0.5p_t & 0.5p_t \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{5}$$

at year t, where its (i, j) element is the transition probability from state i to j, and p_t is given in the formula (3).

In the previous sections we observed a strong effect of the intertemporal elasticity of substitution on the carbon tax. Here, we investigate the joint effect between risk aversion and the uncertain damage level on the carbon tax. We choose the inverse of intertemporal elasticity of substitution as $\psi = 2$, the hazard rate parameter $\nu = 0.00574$. We assume that the tipping point has an uncertain damage level according to equation (4). Figure 6 shows the numbers of carbon tax at the first year for various risk-aversion coefficients and volatilities of the uncertain damage level.

Figure 6: Effect of Risk-Aversion Coefficient and Volatility of Uncertain Damage Level in the Initial Period

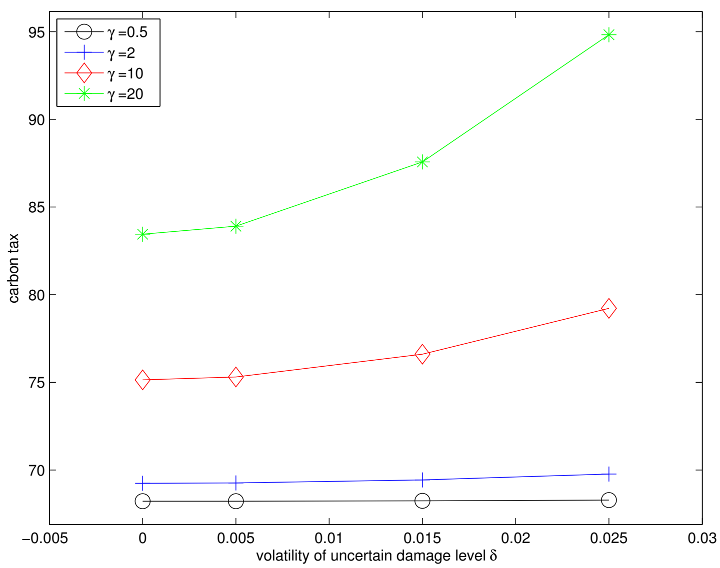

We see that when γ is small, the carbon tax has only a minor increase as the volatility of uncertain damage level increases. However, for larger values of γ, the carbon tax increases in a clearly visible amount as δ increases. Therefore, the effect of a rising uncertainty about future damage from abrupt climate change is amplified with higher degrees of risk aversion. Similar to the basic DICE2007 model with separable preferences it is rather the mean of the uncertain damage and not the variance which affects the carbon tax for low degrees of risk aversion (i.e. smaller than 2). Moreover, the sole effect of the risk-aversion coefficient γ is non-neglectable: it increases much as γ increases from 0.5 to 10 or from 10 to 20, independent of which volatility of the uncertain damage level is chosen, even if the damage level is certain at 5% (i.e., $\delta = 0$).

In the previous example the average damage level was 5% and its volatility at most 2.5%. We now return to the disaster case of the previous section and assume both a higher mean and more uncertain impact. In particular, we assume

$$J_t = \begin{cases} 0.8 + \delta, & \text{with probability } 50\%, \\ 0.8 - \delta, & \text{with probability } 50\%, \end{cases}$$

where $0 \leq \delta < 0.2$, while the probability transition matrix is the same with (5). We choose the inverse of intertemporal elasticity of substitution as $\psi = 2$, and the hazard rate parameter $\nu = 0.0008$, which is the same as in the disaster case of the previous subsection 5.1.3. Again, we study the effect of risk-aversion and volatility of the uncertain damage level on the carbon tax.

Figure 7 shows that for any volatility of the uncertain damage level, the carbon tax significantly increases with higher risk-aversion.

Figure 7: Effect of Risk-Aversion Coefficient and Volatility of Uncertain Damage Level under Disaster Case in the Initial Period

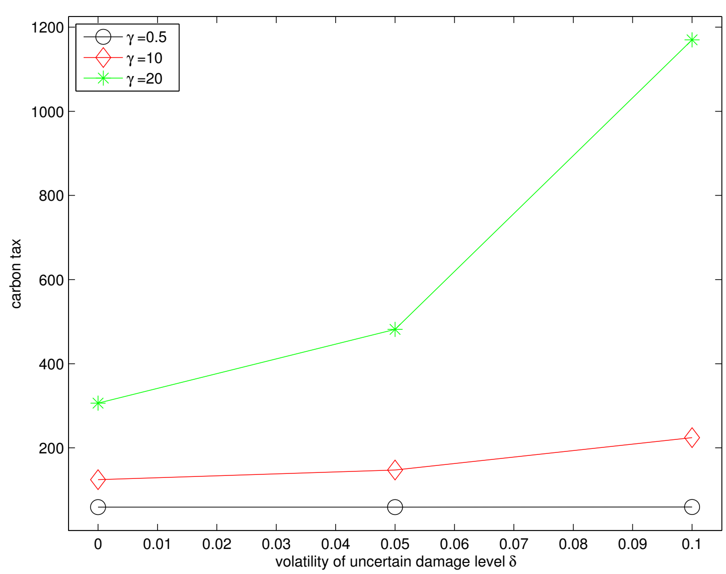

As in the previous example, higher volatility amplifies the effect of risk aversion on the carbon tax. For example, when the volatility is 10%, the carbon tax in the first period is about \$1170 for $\gamma = 20$. This is much higher than the case with $\gamma = 0.5$ (in which the carbon tax is only \$59). Here again, if the risk-aversion coefficient is small (e.g., $\gamma = 0.5$), then higher volatility brings about only a small increase in the carbon tax in the first period. However, if $\gamma = 20$, the carbon tax rises sharply from \$306 to \$1170 when the volatility increases from 0 to 10%. To sum up this example: in face of a very unlikely tipping point event (about 0.1% probability of tipping at 2100), and uncertainty about its impact (mean of 20% and uncertainty is 10%), a risk-averse policymaker sets today's carbon tax at \$1170.

5.3 Multiple Tipping Points

The one-tipping point model illustrates some basic points but is surely not a sensible description of the true range of risks presented by climate change. In this section we look at the effects of multiple tipping events. We focus on two kinds of dynamic tipping processes. First, we examine the case where the events are iid. Second, we look at a situation where there are multiple stages in a compound tipping event. Our example of a multi-stage tipping process can be thought of as representing a sequence of tipping point events. In a recent study Lenton (2012) discusses the domino effect applied to tipping elements, arguing that several tipping elements are sequential and one tipping point might trigger a cascade of additional tipping points. Alternatively, a multi-stage tipping process could represent several consecutive stages of one single tipping element, e.g., sea-level rise due to the melting of GIS. The GIS is a huge ice mass holding an equivalent of about 7m of global sea level (Lenton, 2008). An additional 1 or 2°C of global warming might suffice to trigger an irreversible meltdown of the GIS. Thus, is is quite likely that a GIS tipping point could occur in this century. A tipping of the GIS could lead to a global sea level rise of about 50-100 cm per century (Lenton, 2008). This will enhance the groundling line retreat of the WAIS and cause the sea level to rise even further. In our second multiple-tipping example, we account for a system of three iid tipping elements. Kriegler et al. (2009) estimate the mutual impact of several tipping elements. While there are several interdependent tipping elements, the result suggests a relative independence of the ENSO, the THC and the WAIS. Here again, we do not focus on any specific tipping elements, but rather study in this example how independent tipping elements affects the optimal climate policy.

5.3.1 Three Independent Tipping Points

We first consider the case of three independent tipping points, so we have three discrete state variables, $J_{t,1}$, $J_{t,2}$, $J_{t,3}$, to denote the post-tipping productivities. The i-th tipping point process with damage $J_{t,i}$ follows the probability transition matrix

$$\begin{bmatrix} 1 - p_{t,i} & p_{t,i} \\ 0 & 1 \end{bmatrix},$$

at year t, where $p_{t,i}$ is given in formula (3) with hazard rate parameter ν_i. Therefore, the total number of different 3-dimensional discrete state vector is up to 8. Assuming independence of these three tipping elements, makes this application in essence very similar to the 1 tipping point example with 8 different states. Therefore, the nature of the results for the 3 independent tipping points is very similar to the nature of the previous results.

For simplicity, in this example we assume that each tipping point has the same hazard rate and the same damage. We choose the inverse of intertemporal elasticity of substitution as $\psi = 2$, the risk-averse coefficient $\gamma = 10$, the hazard rate parameter $\nu = 0.01$, and $J_{t,i} = \sqrt[3]{0.95}$ after the tipping event i happens for $i = 1, 2, 3$.

Since now the decision maker faces three independent tipping points simultaneously, the additional carbon tax will incorporate the extra willingness to essentially insure against each tipping point separately. Furthermore, the lower envelope of Figure 8 shows the earliest occurrence of the three tipping points in our simulation runs.

Figure 8: Carbon Tax with Three Tipping Points

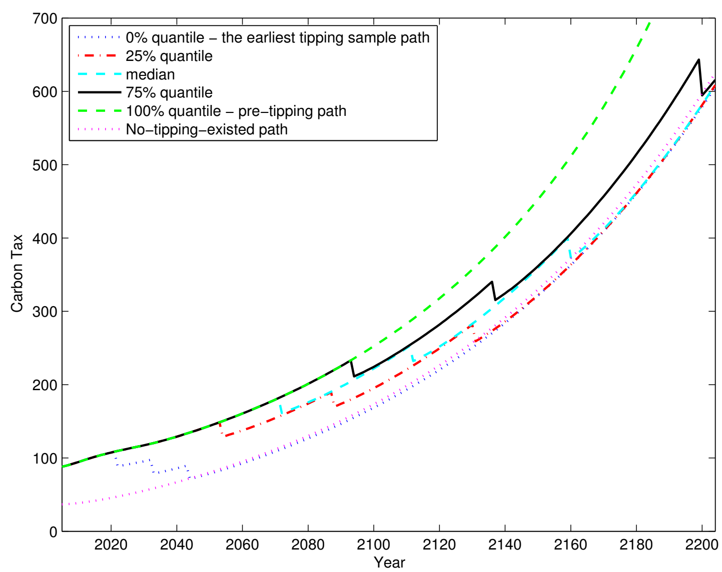

The first tipping point occurs at about 2020, the earliest time at which two out of the three tipping points occur is at 2030, and finally, the earliest time all three tipping point have occurred is 2040. Note that each time one additional tipping point occurred, the carbon tax exhibits the same vertical drop as seen in the 1-tipping point example. Consider now the year 2100. Note, that by 2100, 75% of our simulation runs have exhibited at least one tipping point, 25% of our simulation runs have produced 2 out of the three tipping points, whereas some runs have already produced all three tipping points. Given these characteristics in 2100, the pre-tipping[4] carbon tax is about $250 while the post-tipping tax is about $150. Thus, the optimal preventive carbon tax is around $100.

5.3.2 Four-Stage Tipping Process

We next use DSICE to examine another kind of tipping process. We introduce a four stage tipping process to study how a sequence of tipping events affects the SCC. Since this tipping structure is sequential, the decision maker faces always one tipping point at a time but always has full information about the tipping system transition matrix. To model this, we assume that J_t is a discrete Markov chain with 4 possible values of 1, 59/60, 29/30, 0.95, and its probability transition matrix from year t to year $t+1$ is

$$\begin{bmatrix} 1-p_t & p_t & 0 & 0 \\ 0 & 1-p_t & p_t & 0 \\ 0 & 0 & 1-p_t & p_t \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

where its (i, j) element is the transition probability from state i to j, and p_t is given in the formula (3).

Figure 9 displays the optimal carbon tax for this application, where the inverse of intertemporal elasticity of substitution is $\psi = 2$, the risk-averse coefficient is $\gamma = 10$, and the hazard rate parameter is $\nu = 0.00574$. As in the previous examples, we do report the statistics of 1000 simulated paths of the carbon tax. Given the benchmark parameter specifications, we observe that the carbon tax at the initial period is slightly lower than in the previous examples, despite the same hazard rate. This is because now the decision maker is facing one-third of the full impact at the first stage. Furthermore, the decision maker's expectations of the subsequent tipping stage are now conditional on probability of its preceding stage and the expected timing of the final stage is now much later when compared to the three iid tipping example. The latter point can be seen when comparing the 25% quantile in both plots. According to Figure 9, the 25% quantile for all three tipping events is at around 2170, while in the multi-stage tipping example the 25% quantile for the final stage is much earlier, around 2130. Furthermore, given the sequential nature of the tipping process, the carbon tax exhibits lower volatility before the tipping system reaches its final stage.

[4]For the multiple tipping cases, the pre-tipping scenario implies that no tipping point has occurred yet, while the post-tipping scenario implies that all tipping points under consideration have occurred.

Figure 9: Carbon Tax with a Four-Stage Tipping Process

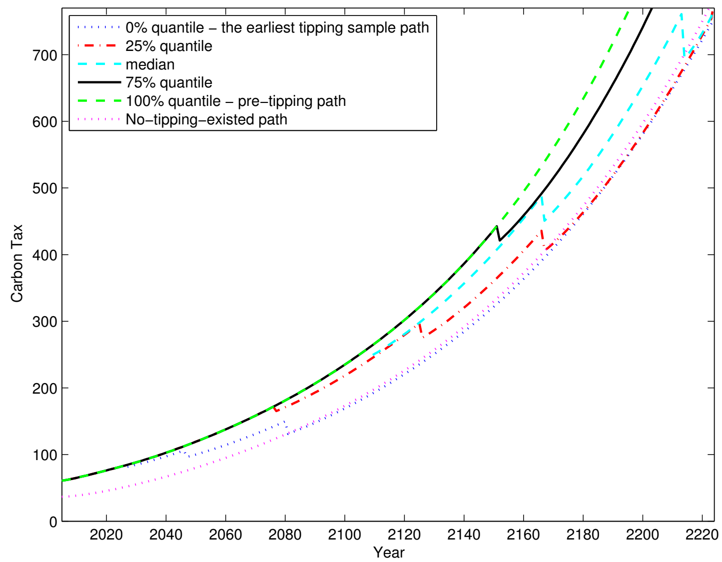

5.4 Shocks to Economic Productivity

We next show how economic productivity shocks interact with climate change dynamics. First, we investigate solely the effect of economic shocks on the carbon tax. We disregard the tipping point for a moment. In the example, we assume that the economic shock ζ_t is a bounded mean-reverting process given in Cai, Judd and Lontzek (2012b): its reverting rate is 0.1, its long-run mean is 1, its long-run standard deviation is 2%, and it is bounded in $(0.92, 1.08)$. We use our baenchmark taste parameters: $\psi = 2$ and $\gamma = 10$, in this example. Figure 10 shows the (relative) difference in the carbon tax to the deterministic model under a long persistence and the benchmark preference parameters. The upper and lower envelopes of our simulations result in roughly a 10% change in the carbon tax. In all simulations, the carbon tax fluctuates between +15% and −15% relative to the mean, with a standard deviation of about 1.7%, a modest amount of volatility.

Figure 10: Effect of Economic Shock

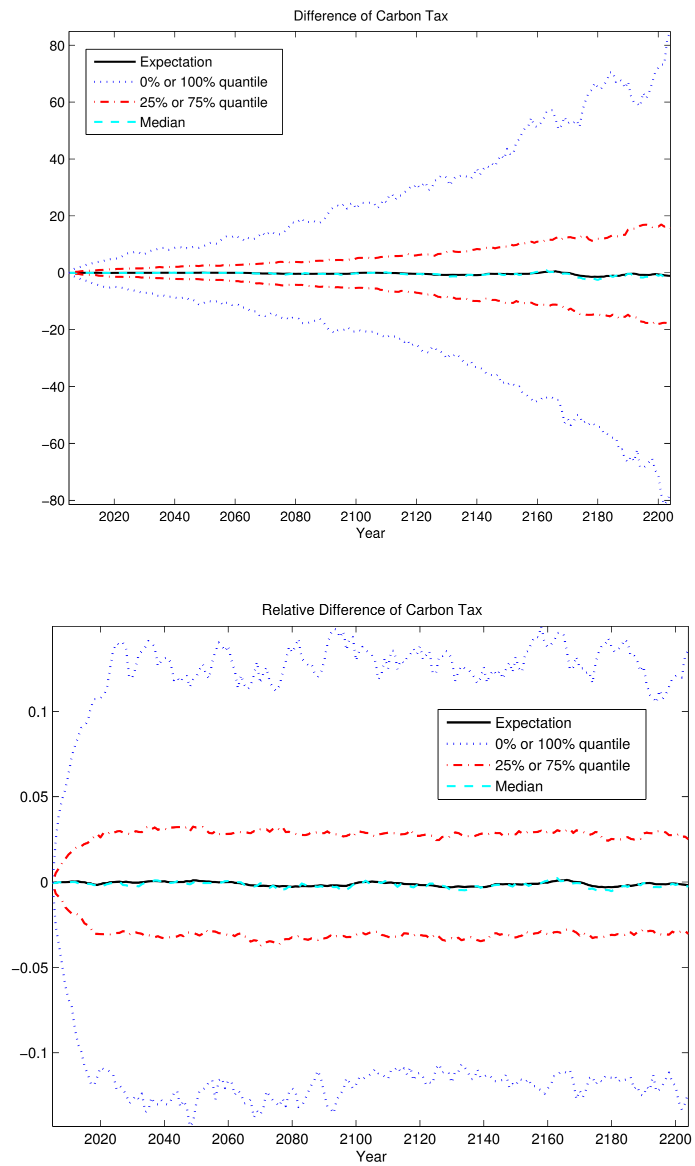

Finally, we present the results of the economic shock uncertainty and the tipping point uncertainty combined. Figure 11 shows the carbon tax for the case with the baseline tipping point with 5% damage level and the business cycle uncertainty having the same parameters with the above example.

Figure 11: Effect of Economic Shock

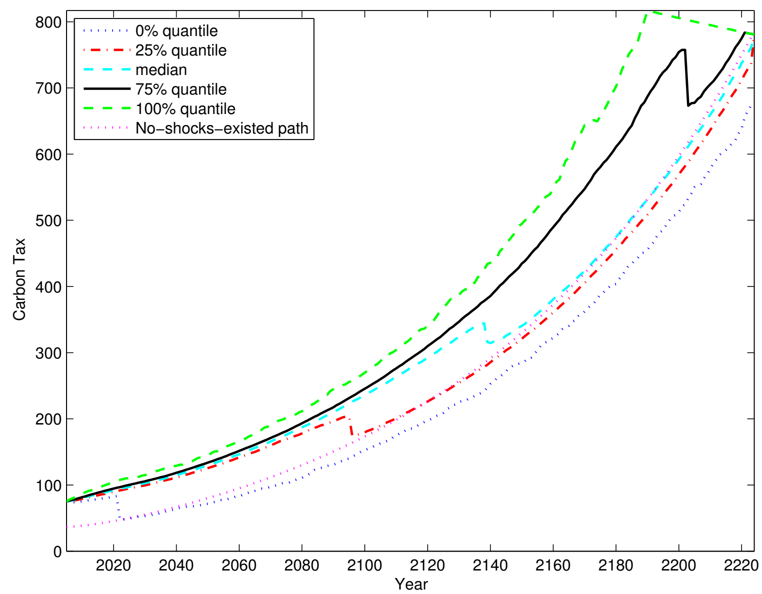

The carbon tax in this example exhibits both features from the "tipping only" and "business cycle only" examples. The constancy in the additional carbon tax from tipping is retained as well as the relatively mild fluctuations of the carbon tax. The latter point can be seen when looking at the positions of the quantiles. As opposed to Figure 2, the optimal carbon taxes for the different quantiles in the post-tipping regime no longer overlap in this example, i.e., they still fluctuates after a tipping point has occurred.

6 Conclusion

This study has analyzed the optimal level and dynamic properties of the carbon tax in face of abrupt and irreversible climate change and its interaction with economic factors, including business cycle fluctuations and preferences about risk. The underlying model entails several ingredients which we consider essential for a sound analysis of climate policy under catastrophic risk. First, we model low-probability tipping point events for abrupt climate change as jump processes with temperature-dependent hazard rates. Second, we account for the intrinsically uncertain nature of future states of the economy and incorporate the cyclicality of business conditions into the optimal decision making process of the policy maker. Third, we model the cost of risk in a manner more compatible with empirical evidence about social risk preferences. Riskiness is inherent in the nature of tipping points, and the socio-economic effects of stochastic, abrupt and irreversible perturbations to the climate system will be affected by preferences about risks. While these features are familiar model components in economics, it is the joint modeling of them within a DSGE version of a widely accepted IAM, that attributes to the novelty of this study.

In particular, we use DSICE (Cai, Judd and Lontzek, 2012b), a DSGE full-dimensional extension of DICE2007. DSICE and DICE2007 are therefore comparable, which facilitates the comparison of our carbon tax numbers with the ones obtained, e.g., by the U.S. Government Interagency Working Group on Social Cost of Carbon (IWG, 2010). In contrast to other approaches in the literature to study a stochastic version of DICE2007, we are endowed with an annual-frequency, full-dimensional, stochastic IAM with intrinsic uncertainty about annual economic productivity and stochastic climate components. We solve DSICE with dynamic programming using advanced computational methods. The solution to DSICE is reliable and quickly obtained. Furthermore, the advanced computational architecture behind DSICE will enable us to study much more complex and higher-dimensional extensions of the model presented in this study.

We find that the threat of a tipping point induces significant and immediate increases in the social cost of carbon, even for low-probability and low-impact tipping events. We find that incorporating Epstein-Zin preferences into DSICE in a manner that is compatible with the evidence on risk aversion and the intertemporal elasticity of substitution significantly increases in the carbon tax. For example, larger values of the intertemporal elasticity of substitution and the degree of risk aversion lead to higher carbon taxes. Furthermore, in addition to abrupt and irreversible climate change, we also study cases with significant uncertainty about post-tipping damages. We find that uncertainty about the damage is also a critical factor leading to a sharp increase in the carbon tax. Furthermore, for low degrees of risk aversion (i.e. smaller than 2), the carbon tax is not affected much by the volatility of the uncertain damage. In contrast, high degrees of risk aversion significantly amplify the effect of damage uncertainty on the carbon tax. We also investigate a disaster scenario that is unlikely (about 0.1% probability of tipping at 2100), but with large and uncertain impacts (mean is 20% and volatility is 10%), today's carbon tax of a highly risk-averse policy maker is $1170.

This application of DSICE also shows that it is quite feasible to solve dynamic stochastic IAMs with current numerical algorithms and computational hardware. Past analyses have been hampered by both software and hardware limitations, but advances over the past twenty years now make it possible to examine important economic questions without making unappealing assumptions for the sake of tractability.

References

[1] Anthoff, D., R.S.J. Tol and G.W. Yohe (2009). Discounting for Climate Change. *Economics: The Open-Access, Open-Assessment E-Journal,* 3(24).

[2] Bahn, O., A. Haurie and R. Malhamé (2008). A Stochastic Control Model for Optimal Timing of Climate Policies, *Automatica*, 44(6), 1545–1558.

[3] Basal, R. & Yaron, A. (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', The Journal of Finance 59(4), 1481–509.

[4] Batjes, J.J. and C.G.M. Goldewijk (1994). The IMAGE 2 hundred year (1890–1990) database of the global environment (HYDE). RIVM, Bilthoven, 410100082.

[5] Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.

[6] Cai, Y. (2009). *Dynamic Programming and Its Application in Economics and Finance*. PhD thesis, Stanford University.

[7] Cai, Y., and K.L. Judd (2010). Stable and efficient computational methods for dynamic programming. *Journal of the European Economic Association*, Vol. 8, No. 2-3, 626–634.

[8] Cai, Y., and K.L. Judd (2012a). Dynamic programming with Hermite interpolation. Working paper.

[9] Cai, Y., and K.L. Judd (2012b). Dynamic programming with shape-preserving rational spline Hermite interpolation. *Economics Letters*, Vol. 117, No. 1, 161–164.

[10] Cai, Y., and K.L. Judd (2012c). Shape-preserving dynamic programming. Forthcoming in *Mathematical Methods of Operations Research*.

[11] Cai, Y., K.L. Judd and T.S. Lontzek (2012a). Continuous-Time Methods for Integrated Assessment Models. RDCEP working paper No. 12-10.

[12] Cai, Y., K.L. Judd and T.S. Lontzek (2012b). DSICE: A Dynamic Stochastic Integrated Model of Climate and Economy. RDCEP working paper No. 12-02.

[13] Cai, Y., K.L. Judd and T.S. Lontzek (2012c). Open science is necessary. *Nature Climate Change*, Vol. 2, Issue 5, 299–299.

[14] Campbell, J. Y. and J. H. Cochrane, (1999), By force of habit: a consumption-based explanation of aggregate stock market behavior, Journal of Political Economy 107, 205–251.

[15] Crost, B. and Traeger, C. P. (2011), Risk and aversion in the integrated assessment of climate change, Department of Agricultural & Resource Economics, UC Berkeley, Working Paper Series 1354288.

[16] Dowlatabati, H., and M.G. Morgan (1993). A Model Framework for Integrated Studies of Climate Change. *Science*, 259, 1813–1814.

[17] Environmental Protection Agency (2010). Peer Review of ADAGE and IGEM.

[18] Epstein, L.G., and S. E. Zin (1989). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57(4), 937–969.

[19] Fisher, A., and U. Narain (2003). Global Warming, Endogenous Risk, and Irreversibility. *Environmental and Resource Economics,* Vol. 25(4), 395–416.

[20] Gill, P., et al. (1994). User's Guide for NPSOL 5.0: a Fortran Package for Nonlinear Programming. Technical report, SOL, Stanford University.

[21] Goettle, R.J., M.S. et al. (2009). Analyzing Environmental Policies with IGEM, an intertemporal General Equilibrium Model of U.S. Growth and the Environment. Cambridge, MA: Dale Jorgenson Associates.

[22] Heal, G., and B. Kriström (2002). Uncertainty and climate change. *Environmental and Resource Economics* 22, 3–39.

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

[24] Interagency Working Group on Social Cost of Carbon (2010). *Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866.* United States Government. `http://www.whitehouse.gov/sites/default/files/omb/inforeg/for-agencies/Social-Cost-of-Carbon-for-RIA.pdf`

[25] Judd, K.L. (1998). *Numerical Methods in Economics.* The MIT Press.

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

[27] Keller, K., B. Bolker and D.F. Bradford (2004). Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management* 48, 723–741.

[28] Kreps, D.M., and E.L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1), 185–200.

[29] Kriegler, I.H., et al. (2009). Imprecise probability assessment of tipping points in the climate system. *PNAS*.

[30] Lemoine, D.M., and C. Traeger (2012). Tipping Points and Ambiguity in the Economics of Climate Change. NBER Working Paper No. 18230.

[31] Lenton, T.M., and J-C. Ciscar (2012). Integrating tipping points into climate impact assessments. *Climatic Change* (in press).

[32] Lontzek, T.S., Y. Cai and K.L. Judd (2012). Tipping points in a dynamic stochastic IAM. RDCEP Working Paper No. 12-03.

[33] Lontzek, T.S., and D. Narita (2011). Risk-averse Mitigation Decisions under an Unpredictable Climate System. *Scandinavian Journal of Economics* 113-4, pp. 937–958.

[34] Mastrandrea, M.D., and S.H. Schneider (2001). Integrated assessment of abrupt climatic changes. *Climate Policy* 1, 433–449.

[35] Manne, A., and R. Richels (2005). MERGE: An Integrated Assessment Model for Global Climate Change. *Energy and Environment* (175-189) edited by. Loulou, R., Waaub, J-P. and Zaccour, G.

[36] McCarl, B., et al. (2011). McCarl GAMS User Guide. GAMS Development Corporation.

[37] Nordhaus, W.D. (1994). *Managing the global commons*. Cambridge, MA: MIT Press.

[38] Nordhaus, W.D. (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.

[39] Nordhaus, W.D. (2011). Estimates of the Social Cost of Carbon: Background and Results from the RICE-2011 Model. Cowles Foundation Discussion Paper No. 1826.

[40] Nordhaus, W.D. and Z. Yang (1996). A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies. *American Economic Review*, 86, 741–765.

[41] Pindyck, R.S. (2008). Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1(1), 46–65.

[42] Pizer, W.A. (1999). The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics* 21, 255–287.

[43] Richardson, L.F., and J.A. Gaunt (1927). The deferred approach to the limit. Philosophical Transactions of the Royal Society of London, Series A 226 (636-646): 299–349.

[44] Rust, J. (2008). Dynamic Programming. In: Durlauf, S.N., Blume L.E. (Eds.), New Palgrave Dictionary of Economics. Palgrave Macmillan, second edition.

[45] Vissing-Jørgensen, A., (2002), Limited asset market participation and the intertemporal elasticity of substitution, Journal of Political Economy 110, 825–853.

[46] Vissing-Jørgensen, A. and O.P. Attanasio (2003), 'Stock-market participation, intertemporal substitution, and risk-aversion', The American Economic Review 93(2), 383–391.

[47] Webster, M., N. Santen and P. Parpas (2012). An Approximate Dynamic Programming Framework for Modeling Global Climate Policy under Decision-Dependent Uncertainty. Computational Management Science.

[48] Wigley, T.M.L., and S.C.B. Raper (1997). Model for the Assessment of Greenhouse-gas Induced Climate Change (MAGICC Version 2.3). The Climate Research Unit, University of East Anglia, UK.

[49] Yohe, G., N. Andronova and M. Schlesinger (2004). To Hedge or Not to Hedge Against an Uncertain Climate Future? *Science* 306, 416–417.

[50] Zickfeld K.et al. (2007). Expert judgements on the response of the Atlantic meridional overturning circulation to climate change. *Climate Change* 82, 235–265.