DSICE: A Dynamic Stochastic Integrated Model of Climate and Economy*

Yongyang Cai
Hoover Institution, 424 Galvez Mall, Stanford University, Stanford, CA 94305
yycai@stanford.edu

Kenneth L. Judd
Hoover Institution, 424 Galvez Mall, Stanford University, Stanford, CA 94305
kennethjudd@mac.com

Thomas S. Lontzek
University of Zurich, Moussonstrasse 15, 8044 Zurich
thomas.lontzek@business.uzh.ch

July 13, 2012

Abstract

This paper introduces a dynamic stochastic integrated model of climate and economy (DSICE), and a numerical dynamic programming algorithm for its solution. More specifically, we solve an example with annual time periods, a six hundred year horizon, and shocks to the economic and climate system. Our dynamic programming methods solve such models on a laptop in about an hour, and do so with good accuracy. This decisively refutes the pessimism one often hears about the possibility of solving such models.

Keywords: numerical dynamic programming, value function iteration, tipping points, stochastic IAM

JEL Classification: C63, Q54, D81

*Cai, Judd, and Lontzek gratefully acknowledge NSF support (SES-0951576). Furthermore, 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.

1 Introduction

There is great uncertainty about the future of the climate and the impact of economic activity on the climate. There is always great uncertainty about future economic conditions. Therefore, any analysis of how society should respond to possible climate change must consider the uncertainties any decision maker faces when choosing policies. This paper demonstrates that there is no difficulty in adding uncertainty and risk to basic IAM models.

We use Nordhaus' DICE2007 model as a starting point for our model. It is well-known in the IAM community and widely used in the IAM literature. Furthermore, it is well-documented. We extend it by adding both economic and climate shocks to the framework of DICE2007.

The IAM community, despite numerous attempts, has so far not been able to produce a stochastic IAM flexible enough to represent uncertainty in a quantitatively realistic manner. In particular, the representation of time should be compatible with the natural frequencies of both the natural and social processes related to climate change. Kelly and Kolstad (1999) is one among only a few models implementing 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. Our approach here is to formulate the general problem and implement it with software and algorithms that can handle a wide rage of problems. Furthermore, we are following standard modern economic notions of uncertainty in terms of time period length and the frequency of shocks, and thus, we use a far more realistic specification of uncertainty. In general, models that assume long time periods, such as ten years, 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. DICE2007 instead uses a ten-year time period. Ten year time periods 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 (2012a) shows that annual time periods produce a significantly different carbon price with the numbers given by DICE2007 (Nordhaus, 2008). Cai, Judd and Lontzek (2012b) develops DICE-CJL, an extension of DICE2007 that can handle any time period length, and then demonstrate that many substantive results depend critically on the time step, strongly supporting our contention that short time periods are necessary for quantitatively reliable analysis.

We then take the DICE-CJL and add economic and climate shocks. We solve it using dynamic programming (DP) methods. In particular, we use

the methods presented in Judd (1998), Cai (2009), and Cai and Judd (2010, 2012a, 2012b, 2012c).

There are many different types of uncertainty that are discussed in the IAM literature. First, many people examine parametric uncertainty, because we do not know with precision the value of key parameters. Second, economic models have substantial amounts of intrinsic uncertainty, meaning that even if one knew the parameters, there would still be uncertainty due to random exogenous events.

DSICE is a model that focuses on intrinsic uncertainty, as is done in the DSGE literature in economics. However, the speed of our DSICE solution algorithm is fast enough that we could also do wide-ranging parameter sweeps to address parameter uncertainty.

We demonstrate the performance of DSICE in a few examples. Those examples show that our algorithm is fast (about an hour on a single-processor laptop to solve a one-year, 600-period stochastic model with eight state variables), and passes basic accuracy tests.

The results show decisively that climate change issues can be examined with the same complexity used in standard dynamic stochastic models in economics.

2 DICE-CJL

DICE2007 (Nordhaus, 2008) maximizes social welfare with tradeoffs between CO_2 abatement, consumption, and investment. In all versions of DICE, the time interval of one period is 10 years, i.e., governments could only have one chance per decade to adjust their economic and climate policy. Moreover, a 10-year time difference is too large in the finite difference method for discretizing the continuous time differential system.

Cai, Judd and Lontzek (2012b) extends DICE2007 to DICE-CJL models that can handle any time period length. Here we summarize and reformulate the annual version of the explicit DICE-CJL model. First, the annual social utility function at year t is

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

where $\gamma = 2$ and c_t is the consumption. We assume that labor supply l_t is inelastic and equals the world population in millions of people, which evolves according to

$$l_t = 6514e^{-0.035t} + 8600(1 - e^{-0.035t}), \tag{1}$$

for any $t = 0, \ldots, 599$. Therefore, the total discounted utility over the first 600 years is

$$\sum_{t=0}^{599} e^{-\rho t} u(c_t, l_t),$$

where $\rho = 0.015$ is the discount rate.

Second, the production side of DICE and DICE-CJL is a basic optimal growth model. Output at year t is produced from capital, k_t (measured in trillions of 2005 U.S. dollars), and labor supply l_t according to the production function

$$f_t(k_t, l_t) = A_t k_t^{\alpha} l_t^{1-\alpha},$$

where $\alpha = 0.3$ is the capital share, and A_t is a total productivity factor defined by

$$A_t = A_0 \exp\left(0.0092(1 - e^{-0.001t})/0.001\right). \tag{2}$$

Global average atmospheric temperature, T_t^{AT} (measured in degrees Celsius above the 1900 temperature), reduces output by a factor

$$\Omega_t = \frac{1}{1 + \pi_1 T_t^{\mathrm{AT}} + \pi_2 (T_t^{\mathrm{AT}})^2},$$

where $\pi_1 = 0$ and $\pi_2 = 0.0028388$. Abatement efforts can reduce CO2 emissions at some cost. Therefore, net output at year t is

$$\mathcal{Y}_t(k_t, T_t^{\mathrm{AT}}, \mu_t) = \frac{1 - \psi_t^{1-\theta_2} \theta_{1,t} \mu^{\theta_2}}{1 + \pi_1 T^{\mathrm{AT}} + \pi_2 (T^{\mathrm{AT}})^2} A_t k^{\alpha} l_t^{1-\alpha},$$

where $\theta_2 = 2.8$, ψ_t is the participation rate assumed to be equal to 1, $\mu_t \in [0, 1]$ is the emission control rate, and

$$\theta_{1,t} = \frac{1.17 \sigma_t \left(1 + e^{-0.005t}\right)}{2\theta_2}, \tag{3}$$

is the adjusted cost for backstop, where σ_t is the technology factor following the path

$$\sigma_t = \sigma_0 \exp\left(-0.0073(1 - e^{-0.003t})/0.003\right).$$

Thus, the next-year capital is

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

where $\delta = 0.1$ is the annual rate of depreciation of capital.

Industrial production processes cause CO2 emissions

$$E_t^{\mathrm{Ind}}(k_t, \mu_t) = \sigma_t (1 - \mu_t) f_t(k_t, l_t),$$

so the total carbon emissions (billions of metric tons) at year t is

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

where

$$E_t^{\mathrm{Land}} = 1.1 e^{-0.01t} \tag{4}$$

is the annual carbon emissions from biological processes.

DICE2007 and DICE-CJL use a simple box model for the carbon cycle. The CO_2 concentrations for the carbon cycle are modeled by a three-layer model,

$$\mathbf{M}_t = (M_t^{\mathrm{AT}}, M_t^{\mathrm{UP}}, M_t^{\mathrm{LO}})^\top,$$

representing carbon concentration (billions of metric tons) in the atmosphere (M_t^{AT}), upper oceans (M_t^{UP}) and lower oceans (M_t^{LO}). The transition system of the CO_2 concentration from year t to year $t+1$ is

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

where

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

where $\phi_{12} = 0.0190837$ and $\phi_{23} = 0.005403$ are calibrated in Cai, Judd and Lontzek (2012b), $\varphi_1 = M_*^{\mathrm{AT}}/M_*^{\mathrm{UP}}$ and $\varphi_2 = M_*^{\mathrm{UP}}/M_*^{\mathrm{LO}}$, where $M_*^{\mathrm{AT}} = 587.473$, $M_*^{\mathrm{UP}} = 1143.894$ and $M_*^{\mathrm{LO}} = 18340$ are the preindustrial equilibrium states of the carbon cycle system.

The CO_2 concentrations impact the surface temperature of the globe through the radiative forcing (watts per square meter from 1900):

$$\mathcal{F}_t\left(M^{\mathrm{AT}}\right) = \eta \log_2\left(M^{\mathrm{AT}}/M_0^{\mathrm{AT}}\right) + F_t^{\mathrm{EX}}. \tag{5}$$

where $\eta = 3.4145$ is calibrated in Cai, Judd and Lontzek (2012b), and F_t^{EX} is the exogenous radiative forcing:

$$F_t^{\mathrm{EX}} = \begin{cases} -0.06 + 0.0036t, & \text{if } t \le 100, \\ 0.3, & \text{otherwise.} \end{cases}$$

DICE2007 and DICE-CJL use a simple box model for the climate. The global mean temperature is represented by a two-layer model,

$$\mathbf{T}_t = (T_t^{\mathrm{AT}}, T_t^{\mathrm{LO}})^\top,$$

representing temperature (measured in degrees Celsius above the 1900 temperature) of the atmosphere (T_t^{AT}) and lower oceans (T_t^{LO}). The transition system of the global mean temperature from year t to year $t+1$ is

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

where

$$\mathbf{\Phi}^{\mathrm{T}} = \begin{bmatrix} 1-\xi_1\eta/3-\xi_1\xi_3 & \xi_1\xi_3 \\ \xi_4 & 1-\xi_4 \end{bmatrix},$$

and $\xi_1 = 0.04217$, $\xi_3 = 0.2609$, and $\xi_4 = 0.0048$ are calibrated in Cai, Judd and Lontzek (2012b).

Therefore, the annual explicit DICE-CJL model is

$$\begin{aligned}
\max_{c_t,\mu_t} \quad & \sum_{t=0}^{599} e^{-\rho t} u(c_t, l_t) + e^{-600\rho} \hat{V}\left(k_{600}, \mathbf{M}_{600}, \mathbf{T}_{600}\right), \\
\text{s.t.} \quad & k_{t+1} = (1-\delta)k_t + \mathcal{Y}_t\left(k_t, T_t^{\mathrm{AT}}, \mu_t\right) - c_t, \\
& \mathbf{M}_{t+1} = \mathbf{\Phi}^{\mathrm{M}} \mathbf{M}_t + \left(\mathcal{E}_t\left(k_t, \mu_t\right), 0, 0\right)^{\top}, \\
& \mathbf{T}_{t+1} = \mathbf{\Phi}^{\mathrm{T}} \mathbf{T}_t + \left(\xi_1 \mathcal{F}_t\left(M_t^{\mathrm{AT}}\right), 0\right)^{\top},
\end{aligned}$$

where $\hat{V}(k, \mathbf{M}, \mathbf{T})$ is the terminal value function defined in Cai, Judd and Lontzek (2012b).

3 DSICE

We add two stochastic shocks to DICE-CJL, one representing an economic shock and the other representing a climate tipping point event that occurs at some random time. The economic shock, denoted by ζ_t, is a continuous stochastic productivity shock corresponding to economic fluctuations 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 a random process. The other novel shock, denoted by J_t, is a jump process that initially equals 1 but then may fall at some future time with the hazard rate of decline related to the contemporaneous temperature, and its transition function from stage t to $t+1$ is $J_{t+1} = g^{J}(J_t, \mathbf{T}_t, \omega_t^{J})$, where ω_t^{J} is another random process independent of ω_t^{ζ} . See Lontzek, Cai, and Judd (2012) for a more complete discussion of the details.

Our expected total discounted utility over 600 years is

$$\mathbb{E}\left\{\sum_{t=0}^{599} e^{-\rho t} u(c_t, l_t)\right\}$$

where $\mathbb{E}\{\cdot\}$ is the expectation operator, and labor supply l_t is defined in the formula (1) of the DICE-CJL model. Our stochastic production function is dependent on both the economic shock and the climate shock:

$$\mathcal{Y}_t(k_t, T_t^{\mathrm{AT}}, \mu_t, \zeta_t, J_t) = \frac{\left(1 - \psi_t^{1-\theta_2} \theta_{1,t} \mu_t^{\theta_2}\right) J_t}{1 + \pi_1 T_t^{\mathrm{AT}} + \pi_2 (T_t^{\mathrm{AT}})^2} \zeta_t A_t k_t^{\alpha} l_t^{1-\alpha}, \tag{6}$$

where the total productivity factor A_t and the adjusted cost for backstop $\theta_{1,t}$ are respectively defined in the formulas (2) and (3) of the DICE-CJL model. Thus, the next-stage capital is

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

The annual total carbon emissions (billions of metric tons) during year t is stochastic and dependent on the economic shock:

$$\mathcal{E}_t\left(k_t, \mu_t, \zeta_t\right) = \sigma_t(1-\mu_t)\zeta_t A_t k_t^{\alpha} l_t^{1-\alpha} + E_t^{\text{Land}}, \tag{7}$$

where E_t^{Land} is the rate of emissions from biological processes defined in the formula (4). Therefore, the carbon cycle and temperature transition system becomes

$$\begin{aligned} \mathbf{M}_{t+1} &= \mathbf{\Phi}^{\text{M}}\mathbf{M}_t + \left(\mathcal{E}_t\left(k_t, \mu_t, \zeta_t\right), 0, 0\right)^{\top}, \\ \mathbf{T}_{t+1} &= \mathbf{\Phi}^{\text{T}}\mathbf{T}_t + \left(\xi_1 \mathcal{F}_t\left(M_t^{\text{AT}}\right), 0\right)^{\top}, \end{aligned}$$

where $\mathcal{F}_t\left(M^{\text{AT}}\right)$ is the annual radiative forcing defined in the formula (5), and $\mathbf{\Phi}^{\text{M}}$ and $\mathbf{\Phi}^{\text{T}}$ are respectively the carbon cycle transition matrix and the climate temperature transfer matrix per year, which are also the same with the annual explicit DICE-CJL model.

Therefore, the stochastic IAM model, called DSICE, becomes

$$\begin{aligned} \max_{c_t, \mu_t} \quad & \mathbb{E}\left\{\sum_{t=0}^{599} e^{-\rho t} u(c_t, l_t) + e^{-600\rho}\hat{V}\left(k_{600}, \mathbf{M}_{600}, \mathbf{T}_{600}, J_{600}\right)\right\} \\ \text{s.t.} \quad & k_{t+1} = (1-\delta)k_t + \mathcal{Y}_t(k_t, T_t^{\text{AT}}, \mu_t, \zeta_t, J_t) - c_t, \\ & \mathbf{M}_{t+1} = \mathbf{\Phi}^{\text{M}}\mathbf{M}_t + \left(\mathcal{E}_t\left(k_t, \mu_t, \zeta_t\right), 0, 0\right)^{\top}, \\ & \mathbf{T}_{t+1} = \mathbf{\Phi}^{\text{T}}\mathbf{T}_t + \left(\xi_1 \mathcal{F}_t\left(M_t^{\text{AT}}\right), 0\right)^{\top}, \\ & \zeta_{t+1} = g^{\zeta}(\zeta_t, \omega_t^{\zeta}), \\ & J_{t+1} = g^{J}(J_t, \mathbf{T}_t, \omega_t^{J}), \end{aligned}$$

where $\hat{V}(k, \mathbf{M}, \mathbf{T}, J)$ is a terminal value function, which will be given in Section 8.

The DP model for DSICE is

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

where V_t is value function at year $t < 600$ and the terminal value function is $\hat{V}$, 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 $\mathbf{M} = (M^{\text{AT}}, M^{\text{UP}}, M^{\text{LO}})^{\top}$ is the three-layer CO_2 concentration and $\mathbf{T} = (T^{\text{AT}}, T^{\text{LO}})^{\top}$ is the two-layer global mean temperature), and $(k^+, \mathbf{M}^+, \mathbf{T}^+, \zeta^+, J^+)$ is its next-year state vector.

4 Numerical Methods for DP

Before discussing examples of DSICE, we summarize the numerical methods we used to solve the dynamic programming problem. In DP problems, if state variables and control variables are continuous such that value functions are also continuous, then we have to use some approximation for the value functions, since computers cannot model the entire space of continuous functions. We focus on using a finitely parameterized collection of functions to approximate value functions, $V(x,\theta) \approx \hat{V}(x,\theta;\mathbf{b})$, 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 (Bellman, 1957). Numerical DP with value function iteration can solve the Bellman equation approximately (Judd, 1998).

A general DP model is based on the Bellman equation:

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

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.

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*

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

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* $\theta_j \in \Theta$.

There are three main components in numerical DP: optimization, approximation, and numerical integration. In the following we focus on discussing approximation and omit the introduction of optimization and numerical integration. Detailed discussion of numerical DP can be found in Cai (2009), Judd (1998) and Rust (2008).

5 Approximation

An approximation scheme consists of two parts: basis functions and approximation nodes. Approximation nodes can be chosen as uniformly spaced nodes, Chebyshev nodes, or some other specified nodes. From the viewpoint of basis functions, approximation methods can be classified as either spectral methods or finite element methods. A spectral method uses globally nonzero basis functions $\phi_j(x)$ such that $\hat{V}(x; \mathbf{c}) = \sum_{j=0}^{n} c_j \phi_j(x)$ is a degree-n approximation. Examples of spectral methods include ordinary polynomial approximation, ordinary Chebyshev polynomial approximation, and shape-preserving Chebyshev polynomial approximation (Cai and Judd, 2012c). In contrast, a finite element method uses local basis functions $\phi_j(x)$ that are nonzero over sub-domains of the approximation domain. Examples of finite element methods include piecewise linear interpolation, shape-preserving rational function spline interpolation (Cai and Judd, 2012b), cubic splines, and B-splines. See Cai (2009), Cai and Judd (2010), and Judd (1998) for more details.

5.1 Chebyshev Polynomial Approximation

Chebyshev polynomials on $[-1, 1]$ are defined as $T_j(z) = \cos(j \cos^{-1}(z))$, while general Chebyshev polynomials on $[a, b]$ are defined as $T_j((2x - a - b)/(b - a))$ for $j = 0, 1, 2, \ldots$. These polynomials are orthogonal under the weighted inner product: $\langle f, g \rangle = \int_a^b f(x)g(x)w(x)dx$ with the weighting function $w(x) = \left(1 - ((2x - a - b)/(b - a))^2\right)^{-1/2}$. A degree n Chebyshev polynomial approximation for $V(x)$ on $[a, b]$ is

$$\hat{V}(x; \mathbf{c}) = \frac{1}{2}c_0 + \sum_{j=1}^{n} c_j T_j \left(\frac{2x - a - b}{b - a} \right), \tag{8}$$

where c_j are the Chebyshev coefficients.

If we choose the Chebyshev nodes on $[a, b]$: $x_i = (z_i + 1)(b - a)/2 + a$ with $z_i = -\cos\left((2i - 1)\pi/(2m)\right)$ for $i = 1, \ldots, m$, and Lagrange data $\{(x_i, v_i) :$

$i = 1, \ldots, m\}$ are given (where $v_i = V(x_i)$), then the coefficients c_j in (8) can be easily computed by the following formula,

$$c_j = \frac{2}{m} \sum_{i=1}^{m} v_i T_j(z_i), \qquad j = 0, \ldots, n. \tag{9}$$

The method is called the Chebyshev regression algorithm in Judd (1998).

It is often more stable to use the expanded Chebyshev polynomial interpolation (Cai and Judd, 2012a), as the above standard Chebyshev polynomial interpolation gives poor approximation in the neighborhood of end points. That is, we use the following formula to approximate $V(x)$,

$$\hat{V}(x; \mathbf{c}) = \frac{1}{2} c_0 + \sum_{j=1}^{n} c_j T_j \left(\frac{2x - \widetilde{a} - \widetilde{b}}{\widetilde{b} - \widetilde{a}} \right), \tag{10}$$

where $\widetilde{a} = a - \delta$ and $\widetilde{b} = b + \delta$ with $\delta = (z_1 + 1)(a - b)/(2z_1)$. Moreover, if we choose the expanded Chebyshev nodes on $[a, b]$: $x_i = (z_i + 1)(\widetilde{b} - \widetilde{a})/2 + \widetilde{a}$, then the coefficients c_j can also be calculated easily by the expanded Chebyshev regression algorithm (Cai, 2009), which is similar to (9).

5.2 Multidimensional Complete Chebyshev Approximation

In a d-dimensional approximation problem, let the domain of the approximation function be

$$\{x = (x_1, \ldots, x_d) : a_i \le x_i \le b_i, \, i = 1, \ldots d\},$$

for some real numbers a_i and b_i with $b_i > a_i$ for $i = 1, \ldots, d$. Let $a = (a_1, \ldots, a_d)$ and $b = (b_1, \ldots, b_d)$. Then we denote $[a, b]$ as the domain. Let $\alpha = (\alpha_1, \ldots, \alpha_d)$ be a vector of nonnegative integers. Let $T_\alpha(z)$ denote the product $T_{\alpha_1}(z_1) \cdots T_{\alpha_d}(z_d)$ for $z = (z_1, \ldots, z_d) \in [-1, 1]^d$. Let

$$Z(x) = \left(\frac{2x_1 - a_1 - b_1}{b_1 - a_1}, \ldots, \frac{2x_d - a_d - b_d}{b_d - a_d} \right)$$

for any $x = (x_1, \ldots, x_d) \in [a, b]$.

Using these notations, the degree-n complete Chebyshev approximation for $V(x)$ is

$$\hat{V}_n(x; \mathbf{c}) = \sum_{0 \le |\alpha| \le n} c_\alpha T_\alpha \left(Z(x) \right),$$

where $|\alpha| = \sum_{i=1}^{d} \alpha_i$ for the nonnegative integer vector $\alpha = (\alpha_1, \ldots, \alpha_d)$. So the number of terms with $0 \le |\alpha| = \sum_{i=1}^{d} \alpha_i \le n$ is $\binom{n+d}{d}$ for the degree-n complete Chebyshev approximation in $\mathbb{R}^d$.

Let

$$z^{(k)} = \left(z_1^{(k_1)}, \ldots, z_d^{(k_d)} \right) \in [-1, 1]^d,$$

where $k = (k_1, \ldots, k_d)$, $z_i^{(k_i)} = -\cos\left((2k_i - 1)\pi/(2m)\right)$ for $k_i = 1, \ldots, m$, and $i = 1, \ldots, d$. Let

$$x_i^{(k_i)} = (z_i^{(k_i)} + 1)(b_i - a_i)/2 + a_i,$$

for $i = 1, \ldots, d$, and then

$$x^{(k)} = \left(x_1^{(k_1)}, \ldots, x_d^{(k_d)}\right)$$

is a d-dimensional Chebyshev node in $[a, b]$. These $x^{(k)}$ (for all $k_i = 1, \ldots, m$ and $i = 1, \ldots, d$) forms the set of the d-dimensional Chebyshev nodes with m nodes in each dimension. For each $x^{(k)}$, let $v^{(k)} = V(x^{(k)})$ be computed by solving the Bellman equation at $x^{(k)}$. Then the coefficients of the degree-n complete Chebyshev approximation on $[a, b]$ are computed as

$$c_\alpha = \frac{2^{\tilde{d}}}{m^d} \sum_{1 \le k_i \le m, 1 \le i \le d} v_{(k)} T_\alpha(z^{(k)}),$$

where $\tilde{d} = \sum_{i=1}^{d} 1_{\alpha_i > 0}$ with $1_{\alpha_i > 0}$ as the indicator

$$1_{\alpha_i > 0} = \begin{cases} 1, & \text{if } \alpha_i > 0, \\ 0, & \text{if } \alpha_i = 0, \end{cases}$$

for all nonnegative integer vectors α with $0 \le |\alpha| \le n$.

We can easily extend this multidimensional complete Chebyshev approximation and the formula to compute the Chebyshev coefficients to its expanded version over $[a, b]$.

6 Discrete-Time Bounded Process of Continuous Economic Shock

An economic shock process usually has a mean-reverting property: when it is higher than its long-run mean, it tends to drift down; when it is lower than its long-run mean, it tends to drift up; and the tendency will be higher if it is more away from its long-run mean. A simple discrete-time AR(1) mean-reverting process y_t is:

$$y_{t+1} = (1 - \lambda) y_t + \sqrt{2\lambda - \lambda^2} z_t,$$

where $0 < \lambda < 1$ is the rate of reversion and z_t is a normal random variable with zero mean and unit variance, and the z_t process is i.i.d. Therefore, y_t is a normal random variable with long-run zero mean and long-run unit variance, and unbounded.

Figure 1: Difference between standard normal random variable y and its transformed bounded random variable ζ

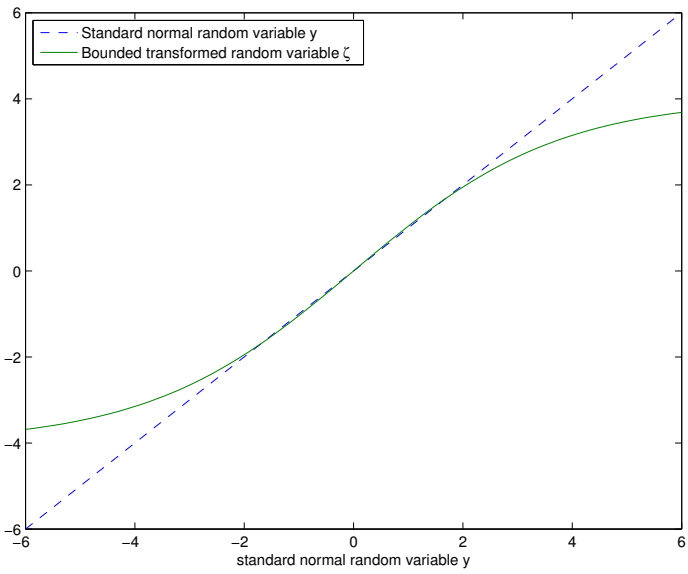

However, an economic shock should be bounded. We will use the following function to transform the standard normal random variable y to a bounded random variable ζ:

$$\zeta = \frac{1 - e^{-\kappa y}}{1 + e^{-\kappa y}} \nu, \tag{11}$$

where ν and κ are two parameters. We see that ζ is bounded in $(-\nu, \nu)$ and its mean is zero as the transformation function is an odd function of y. Once we choose a number of ν, we would like to choose a corresponding κ so that the variance of ζ is one. For example, if we set $\nu = 4$, then we will choose $\kappa = 0.532708$.

Figure 1 shows the difference between the standard normal random variable and its transformed bounded random variable using the formula (11) with $\nu = 4$ and $\kappa = 0.532708$. We see that when the standard normal random variable y has a value in $(-2, 2)$ having up to 95% probability, its transformed bounded random variable ζ is close to y. Figure 2 shows the probability density function of ζ which tells us that it is bounded within $(-4, 4)$ and symmetric at its mean 0.

Now we can generate a discrete-time bounded mean-reverting process for the economic shock ζ_t using the following function:

$$\zeta_t = \bar{\zeta} + \frac{1 - e^{-\kappa y}}{1 + e^{-\kappa y}} \nu\sigma, \tag{12}$$

Figure 2: Probability density functions of standard normal random variable y and its transformed bounded random variable ζ

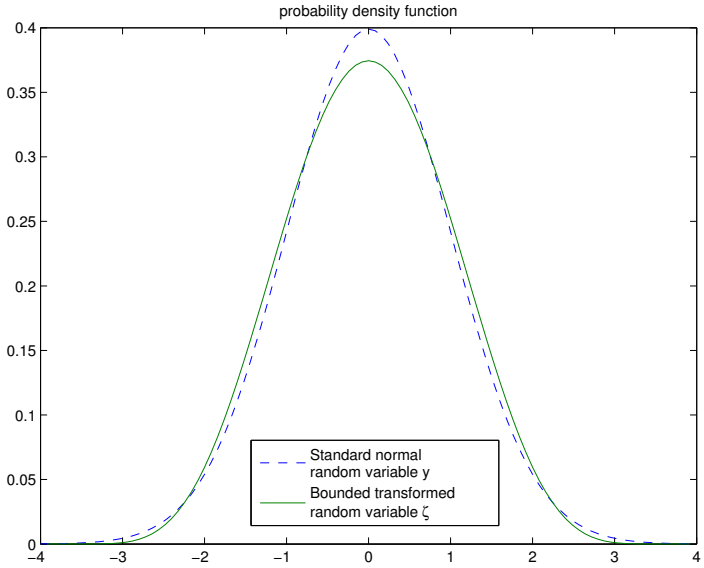

so that its long-run mean is $\bar{\zeta}$, its reverting rate is λ, its long-run variance is σ^2, and it is bounded in $\left(\bar{\zeta} - \nu\sigma, \bar{\zeta} + \nu\sigma\right)$. In our examples, we will have $\bar{\zeta} = 1$, and we should have $\bar{\zeta} - \nu\sigma > 0$ to guarantee that the economic shock ζ_t is always positive.

7 Choosing Domains of Value Functions

In our examples, $(k, \mathbf{M}, \mathbf{T}, \zeta)$ is the 7-dimensional vector of continuous states, and J is one discrete state. All of our examples assume that we solve a 600-year horizon problem. As is assumed in DICE, we believe that this solution of the first few centuries are a good approximation of the solution of the first few centuries to the infinite horizon problem in the model description.

It is important to set an appropriate domain for approximating the value functions. We use the solution of DICE-CJL to tell us how to construct the domain of the value function at each time. Over these 600 years, the optimal solution of the annual DICE-CJL tells us that the minimal capital is 137, and the maximal capital is 70138 along the optimal path of capital. Therefore, if we use a fixed domain along the time path, then the domain will be too large. The problem becomes more difficult when we include the stochastic states, particularly the tipping point shock, in DSICE.

To overcome this difficult problem, we let the domains vary along the time path. We use the optimal solution of DICE-CJL to generate the domains

of the value functions along years, and keep the optimal state variables of DICE-CJL at around the center of the domain. When there is a tipping point shock in the model, the optimal state variables of the adjusted DICE-CJL should also be around the center of the domain, where the adjusted DICE-CJL assume that the tipping point happens at the first year.

To choose the domains with the above properties, we set the range of capital k_t along year t to be

$$\left[0.75 \min_{\zeta,J} \left\{\zeta k^*_{t,J}\right\}, 1.2 \max_{\zeta,J} \left\{\zeta k^*_{t,J}\right\}\right], \tag{13}$$

where $k^*_{t,J}$ is the optimal capital path of adjusted DICE-CJL with the production function $\mathcal{Y}_t(k_t, T^{\mathrm{AT}}_t, \mu_t, 1, J)$ for a fixed tipping level J over the 600-year horizon. The ranges of the other continuous states are defined in the following way:

$$\begin{aligned}
\underline{\mathbf{M}}_{t+1} &= \mathbf{\Phi}^{\mathrm{M}} \underline{\mathbf{M}}_t + (\underline{E}_t, 0, 0)^\top, \\
\overline{\mathbf{M}}_{t+1} &= \mathbf{\Phi}^{\mathrm{M}} \overline{\mathbf{M}}_t + (\overline{E}_t, 0, 0)^\top, \\
\underline{\mathbf{T}}_{t+1} &= \mathbf{\Phi}^{\mathrm{T}} \underline{\mathbf{T}}_t + (\xi_1 \underline{F}_t, 0)^\top, \\
\overline{\mathbf{T}}_{t+1} &= \mathbf{\Phi}^{\mathrm{T}} \overline{\mathbf{T}}_t + (\xi_1 \overline{F}_t, 0)^\top,
\end{aligned}$$

where $\underline{\mathbf{M}}_t$ and $\underline{\mathbf{T}}_t$ are the lower bounds of $\mathbf{M}_t$ and $\mathbf{T}_t$ respectively, $\overline{\mathbf{M}}_t$ and $\overline{\mathbf{T}}_t$ are the upper bounds of $\mathbf{M}_t$ and $\mathbf{T}_t$ respectively,

$$\begin{aligned}
\underline{F}_t &= \eta \log_2 \left(\underline{M}^{\mathrm{AT}}_t / M^{\mathrm{AT}}_0\right) + F^{\mathrm{EX}}_t, \\
\overline{F}_t &= \eta \log_2 \left(\overline{M}^{\mathrm{AT}}_t / M^{\mathrm{AT}}_0\right) + F^{\mathrm{EX}}_t,
\end{aligned}$$

and $\underline{E}_t$ and $\overline{E}_t$ are the lower and the upper bounds of the emission, $\mathcal{E}_t\left(k_t, \mu_t, \zeta_t\right)$, at year t.

8 Terminal Value Function in DSICE

Assume that at the terminal time, the capital is $\tilde{k}$, the three-layer CO_2 concentration is $\tilde{\mathbf{M}}$, the two-layer global mean temperature is $\tilde{\mathbf{T}}$. We assume that $h = 1$ for the dynamic system after the terminal time (the 600th year). For any time after the terminal time, we assume that the population is $l_t = 8600$, the total production factor and the adjusted cost for backstop will be the same with the numbers at the terminal time respectively, i.e., $A_t = 1.7283$ and $\theta_{1,t} = 0.00386$. DSICE assumes that at the terminal time, the world reaches a partial equilibrium: after the terminal time, capital will be the same, and the emission control rate will always be 1 so that the emission of carbon from the industry will always be 0, i.e., $k_t = \tilde{k}$ and $\mu_t = 1$, for any year $t \geq 600$. Moreover, the economic shock disappears,

and the tipping point event will never happen after the terminal time, i.e., $\zeta_t = 1$ and $J_t = \tilde{J}$ where $\tilde{J}$ is the tipping level at the terminal time. Thus, the dynamics of the climate system become

$$\begin{aligned} \mathbf{M}_{t+1} &= \mathbf{\Phi}^{\mathrm{M}}\mathbf{M}_t + (E_t^{\mathrm{Land}}, 0, 0)^\top, \\ \mathbf{T}_{t+1} &= \mathbf{\Phi}^{\mathrm{T}}\mathbf{T}_t + (\xi_1 F_t, 0)^\top, \end{aligned}$$

for any year $t \geq 600$, where $\mathbf{M}_{600} = \tilde{\mathbf{M}}$, $\mathbf{T}_{600} = \tilde{\mathbf{T}}$, and

$$F_t = \eta \log_2 \left(M_t^{\mathrm{AT}}/M_0^{\mathrm{AT}}\right) + 0.3.$$

To keep the above partial equilibrium, the consumption is

$$c_t = \mathcal{Y}_t(k_t, T_t^{\mathrm{AT}}, 1, 1, \tilde{J}) - \delta k_t.$$

Therefore, we have our terminal value function:

$$V(\tilde{k}, \tilde{\mathbf{M}}, \tilde{\mathbf{T}}, \tilde{J}) = \sum_{t=600}^{\infty} e^{-\rho(t-600)} u(c_t, l_t).$$

To compute the terminal value function, we will use the summation of discounted utilities over 800 years from $t = 600$ to $t = 1399$ with one year as the time interval for each period instead. It will be a very good approximation of the summation of the infinite sequence, because $e^{-800\rho} \approx 6.1 \times 10^{-6}$ is small enough. That is,

$$V(\tilde{k}, \tilde{\mathbf{M}}, \tilde{\mathbf{T}}, \tilde{J}) \approx \sum_{t=600}^{1399} e^{-\rho(t-600)} u(c_t, l_t).$$

It would be too time-consuming to use the terminal value function of the above formula in optimizers to compute optimal solutions, so we will use its approximation to save computational time. In our examples, for each possible value of J, we will use a complete Chebyshev polynomial approximation, $\hat{V}(k, \mathbf{M}, \mathbf{T}, J)$, over the terminal domain of the 6-dimensional continuous state space of $(k, \mathbf{M}, \mathbf{T})$, which is chosen by the approach discussed in Section 7.

9 Accuracy Test

An accuracy test is very important for any numerical algorithm. A numerical algorithm should not be trusted if we have not examined its accuracy. Here again the DICE-CJL model helps us. It will be a natural way to compare the solutions given by the numerical DP algorithm and the solutions given by the GAMS code using a large-scale nonlinear optimizer for DICE-CJL in Cai, Judd and Lontzek (2012b), because DICE-CJL are degenerated cases of DSICE.

Our Fortran code of the numerical DP algorithm (the deterministic version of Algorithm 1 with six continuous state variables and a degenerated economic shock) is applied to solve DICE-CJL. In the maximization step of DP, we use NPSOL (Gill et al., 1994), a set of Fortran subroutines for minimizing a smooth function subject to linear and nonlinear constraints. For each dimension of the continuous state space, we choose 5 expanded Chebyshev nodes, and then use the tensor rule to generate all points. In our examples, the number of points is $5^6 = 15,625$. Therefore, for each value function iteration, we compute 15,625 values of the value function at these points, and then compute Chebyshev coefficients of a degree-4 expanded complete Chebyshev polynomial to approximate the value function.

After computing the Chebyshev coefficients for all stages along the 600 years using the backward value function iteration method, we generate the optimal path with the given initial state in the GAMS code by the forward iteration method. That is, given the current stage's state, since we have the approximation of the next-stage value function, we can use the Bellman equation to compute the optimal consumption and emission control so that we can get the optimal next-stage state, and then go on until the terminal stage.

Now we use the solution of DICE-CJL from the GAMS code in Cai, Judd and Lontzek (2012b) to verify the accuracy of the optimal path computed from the numerical DP algorithm. Table 1 lists the relative errors of the optimal paths of states, $(k_t, M_t^{\mathrm{AT}}, T_t^{\mathrm{AT}})$, and control variables, (c_t, μ_t), over the first 400 years and the running time of the numerical DP algorithm, which are run on a laptop. The relative errors of the other states, $(M^{\mathrm{UP}}, M^{\mathrm{LO}}, T^{\mathrm{LO}})$, are even smaller so they are omitted in the table. The errors are computed in the following formula:

$$\max_{t \le 400} \left| \frac{x^*_{t,\mathrm{DP}} - x^*_{t,\mathrm{GAMS}}}{x^*_{t,\mathrm{GAMS}}} \right|,$$

where $x^*_{t,\mathrm{DP}}$ is the optimal path at year t from our numerical DP algorithm, and $x^*_{t,\mathrm{GAMS}}$ is the optimal solution at year t of the GAMS code for the annual explicit DICE-CJL model. When we use the degree-4 complete Chebyshev polynomial as the approximation method in the DP algorithm, the relative errors are small, varying from $O(10^{-4})$ to $O(10^{-5})$. When we increase the degree to 6, the accuracy is up to $O(10^{-5})$ or $O(10^{-7})$.

However, the running time of numerical DP with the degree-4 complete Chebyshev polynomial is only 8.5 minutes, while the running time of numerical DP with the degree-6 complete Chebyshev polynomial is up to 2.7 hours, about 19 times more time, on the same laptop. The main reason of this big difference of running times is that we need to compute values at $7^6 = 117,649$ points for degree-6 complete Chebyshev polynomial approximation while the degree-4 one has only $5^6 = 15,625$ points, so the degree-6 one needs to com-

Table 1: Relative Errors and Running Time of the Numerical DP Algorithm

degree	k	M^{AT}	T^{AT}	c	μ	Time
4	6.4(−4)	5.7(−5)	7.2(−5)	2.0(−4)	8.5(−5)	8.5 minutes
6	6.6(−6)	6.2(−7)	4.5(−7)	1.7(−5)	2.0(−6)	2.7 hours

Note: $a(-n)$ means $a \times 10^{-n}$.

pute about 7.5 times more points than the degree-4 one. The second reason is that the degree-6 complete polynomial has $\binom{6+6}{6} = 924$ terms, while the degree-4 complete polynomial has only $\binom{4+6}{6} = 210$ terms, so that in the objective function of the maximization step of Algorithm 1, the running time of next-stage value function approximation in degree-6 complete polynomial takes about 4 times more time than the degree-4 one.

The numerical DP algorithm with the degree-6 complete polynomial is much more time-consuming than the one with the degree-4 one to solve DICE-CJL. Furthermore, the degree-4 one has enough accuracy. Therefore, we will keep using the degree-4 complete Chebyshev polynomial approximation in the numerical DP algorithm in the stochastic examples shown later. At this state we would like to point out, that the prime objective of this study is to demonstrate the ability of our numerical algorithm to deal with higher-dimensional stochastic problems. We abstain here from an in-depth analysis of, e.g., the carbon tax for which one would need to analyze alternative calibrations and sensitivity to different utility functions.

10 DSICE Results

In this section we apply numerical DP algorithms in our Fortran code to solve DSICE. For each discrete state value, we choose the degree-4 expanded complete Chebyshev polynomials to approximate the value functions. Moreover, we compute the values of the value function on the multidimensional tensor grids with 5 expanded Chebyshev nodes on each dimension in the continuous state ranges, and then compute the Chebyshev coefficients.

After computing these Chebyshev coefficients on all discrete state values for all stages along the 600 years using the backward value function iteration method, we use a simulation method to generate the optimal paths by the forward iteration method. That is, when the state at the current stage is given, since the next stage value function approximation has been computed by the previous numerical DP algorithm, we can apply the optimization solver to get the optimal policy and the next-stage continuous state $(k, \mathbf{M}, \mathbf{T})$. Then we simulate to get the next stage stochastic state. We start this process with the given initial continuous state in the GAMS code and $(\zeta_0, J_0) = (1, 1)$, and run it until the terminal time. In the following examples, we will compute 1000 optimal paths by simulation method, and

then plot their distribution.

10.1 DSICE with an Economic Shock

In this example, we consider the economic shock ζ_t in the model. Let ζ_t be a discrete-time bounded mean-reverting process with long-run mean $\bar{\zeta} = 1$, reverting rate $\lambda = 0.1$, and long-run standard deviation $\sigma = 2\%$.

Figures 3-6 show the numerical results of DSICE. Figure 3 shows the distribution of optimal paths of capital over the first 200 years. The solid line is the average optimal capital along time t, the dotted lines are the minimal or maximal optimal capital, the dashed line is the median optimal capital, the dash-dot lines are the 25% or 75% quantile of the 1000 optimal paths, and the marked lines are the lower bounds or the upper bounds of the range of capital for the value function approximation which are given by the formula (13). Figures 4, 5 and 6 plot the distribution of optimal paths of carbon concentration in the atmosphere, surface temperature and carbon tax respectively. From Figures 3, 4 and 5, we see that all of 1000 simulated paths are located inside and around the center of the corresponding wide approximation ranges for capital, atmospheric carbon and surface temperature (this is also true for the other three state variables). This tells us that our optimal solutions of capital are not binding on the bounds of the approximation range and then the value function approximation is smooth without kinks so that it enables the reliable implementation of the numerical DP algorithm (Algorithm 1).

We see that the economic shock has a significant impact on both the capital stock and the carbon tax, but little impact on either carbon concentration or temperature. This happens because in this example the economic shock is not irreversible, in this example. Moreover, when we have more money from the economic shock with $\zeta_t > 1$, we will spend more money to reduce CO_2 emissions, and when we have less money with $\zeta_t < 1$, we will spend less money in reducing CO_2 emissions.

10.2 DSICE with a Tipping Point

In this example, we consider a tipping point shock J_t in DSICE. Let J_t be a discrete Markov chain with 2 possible values of $1.0, 0.9$, and its 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

$$p_t = \max\left\{0, \min\left\{1, \ \frac{h(T_t^{\mathrm{AT}} - 1)}{100}\right\}\right\},$$

where h is the length of time step size in years (in this example, $h = 1$ year). So the probability p_t is dependent on the surface temperature at time t, higher surface temperature implies higher probability to have an irreversible damage, as the transition probability of J_t from state 2 to 1 is 0. See Lontzek, Cai and Judd (2012) for a more detailed discussion about DSICE with tipping points.

Figure 7 shows the distribution of optimal paths of capital over the first 200 years. The solid line is the average optimal capital along time t, the dotted lines are the minimal or maximal optimal capital, the dashed line is the median optimal capital, the dash-dot lines are the 25% or 75% quantile of the 1000 optimal paths, and the marked lines are the lower bounds or the upper bounds of the range of capital for the value function approximation which are given by the formula (13). Figures 8, 9 and 10 plot the distribution of optimal paths of carbon concentration in atmosphere, surface temperature and carbon tax respectively. From Figures 7, 8 and 9, we see that all of 1000 simulated paths are located inside and around the center of the corresponding wide approximation ranges for capital, atmospheric carbon and surface temperature (this is also true for the other three state variables). This tells us that our optimal solutions of capital are not binding on the bounds of the approximation range and then the value function approximation is smooth without kinks so that it enables the reliable implementation of the numerical DP algorithm (Algorithm 1).

From these figures, we see that among these 1000 optimal paths, the tipping point has a significant impact on all of these state variables and control variables. Once the tipping event happens, carbon concentration and temperature increase dramatically, the emission control rate exhibits a big jump, and capital will stop growing in the first years and then start growing at a much smaller speed.

10.3 DSICE with an Economic Shock and a Tipping Point

Our last example considers the combination of an economic shock and a tipping point shock in the above two examples. Figure 11 shows the distribution of optimal paths of capital over the first 200 years. The solid line is the average optimal capital along time t, the dotted lines are the minimal or maximal optimal capital, the dashed line is the median optimal capital, the dash-dot lines are the 25% or 75% quantile of the 1000 optimal paths, and the marked lines are the lower bounds or the upper bounds of the range of capital for the value function approximation which are given by the formula (13). Figure 12 plots the distribution of optimal paths of the carbon tax. The distribution of carbon concentration in the atmosphere and the surface temperature are similar to the previous example of DSICE with a tipping point. From Figure11, we see that all of 1000 simulated paths are located inside and around the center of the corresponding wide approximation ranges

Table 2: Running Time of the Numerical DP Algorithm for DSICE

	Time
DSICE with an Economic Shock	2.0 hours
DSICE with a Tipping Point	17.5 minutes
DSICE with an Economic Shock and a Tipping Point	4.1 hours

for capital (this is also true for the other five state variables). This tells us that our optimal solutions of capital are not binding on the bounds of the approximation range and then the value function approximation is smooth without kinks so that it enables the reliable implementation of the numerical DP algorithm (Algorithm 1).

From the figures, we see that a tipping point has a significant impact on the state and control variables, and the economic shock also affects the capital and carbon tax significantly. Particularly for the carbon tax (Figure 12), we see that the dotted line with maximal carbon tax is not at the same path with the dash-dot or dashed lines before the tipping event happens, but they are the same in Figure 10. This difference is caused by the impact of economic shock.

10.4 Running Times

Table 2 lists the running times of the numerical DP algorithm with 600 value function iterations for solving the previous examples of DSICE, which are run on a single-core laptop. We see that the problem of DSICE with a tipping point takes only 17.5 minutes, while the maximal time is also only 4.1 hours for solving the problem of DSICE with both an economic shock and a tipping point. The reason that the problem with an economic shock takes much more time than the one with only a tipping point is that the economic shock is a continuous state variable and the tipping point is a discrete state variable with only two values. The small running times in Table 2 tell us that the numerical DP algorithm is fast in solving the stochastic IAMs even with one year as the time step.

10.5 Longer Step Size

Climate and economy are both continuous-time systems. The mutual interaction between these two systems forms the core of any Integrated Assessment Model. Nevertheless, it is common practice in the IAM literature to specify the climate-economy in discrete time, typically assuming very long discrete time-steps of 5 or 10 years (or even much longer in many stochastic IAMs). If IAMs use e.g. decadal time steps, it would be highly desirable that they properly represent the true continuous time dynamics of the underlying system and address the appropriate policies to cope with adverse

effects of climate change. The insights obtained from IAMs are frequently used by policy makers to design and evaluate various climate policies, such as carbon taxes and global warming targets. For example the United States government (Interagency Working Group on Social Cost of Carbon, 2010) has recently engaged in determining the social costs of carbon, the dollar value on damages from one more ton of carbon emissions. The DICE model was one the three models used for this analysis.

However, as we show in Cai, Judd and Lontzek (2012a, 2012b), the deterministic DICE or DICE-CJL model with a step size longer than one year gives an irreliable solution. Here we also show that if we choose a long step size of 5 or 10 years for the DSICE model, then the optimal solutions are not reliable too, even when one uses the best numerical methods like Algorithm 1.

Figure 13 shows that the mean of the optimal carbon tax path over the 1000 simulation paths for DSICE with 1-year, 5-year and 10-year time step sizes. We see that the 1-year carbon tax is always a lot smaller than 5-year or 10-year numbers. In particular, at year 2015, the optimal carbon tax of 1-year step size is \$142, but the number of 5-year step size is \$159 which is 12% higher, and the number of 10-year step size is \$171 which is 20% higher than the one of 1-year step size.

11 Conclusion

We have described DSICE, a basic dynamic stochastic extension of DICE2007. We have shown that it is quite feasible to combine annual (even sub-annual) time periods with economic and climate uncertainty. The speed of the algorithms means that we can do an extensive exploration of the parameter space to determine the sensitivity of conclusions for parameters about which we have limited information. The accuracy tests indicate that the algorithms are reliable as well as fast. This paper has focused on describing the basic model and addressing basic issues relating to the feasibility of such a model. We have clearly refuted the pessimism one often hears about the possibility of such analyses.

References

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

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

[3] 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.

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

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

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

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

[8] Cai, Y., K.L. Judd and T.S. Lontzek (2012b). Continuous-Time Methods for Integrated Assessment Models. Working paper.

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

[10] 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`

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

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

[13] Lontzek, T.S., Y. Cai and K.L. Judd (2012). Tipping points in a dynamic stochastic IAM. Working paper.

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

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

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

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

Figure 3: Capital in DSICE with an Economic Shock

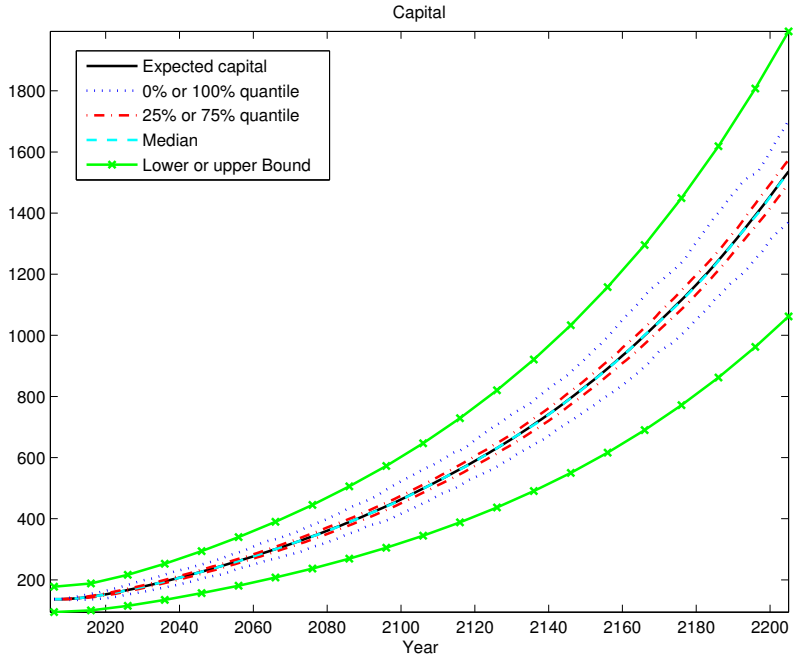

Figure 4: Atmospheric Carbon in DSICE with an Economic Shock

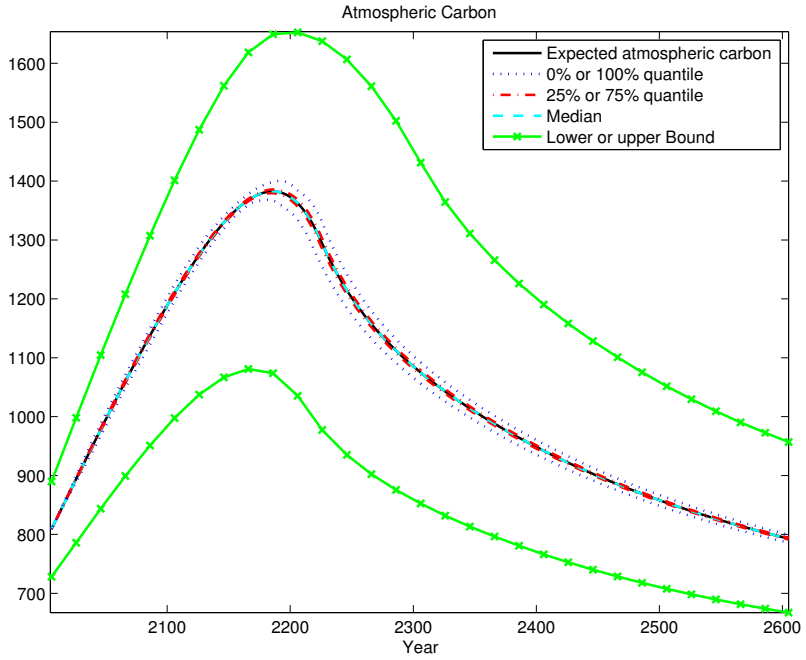

Figure 5: Surface Temperature in DSICE with an Economic Shock

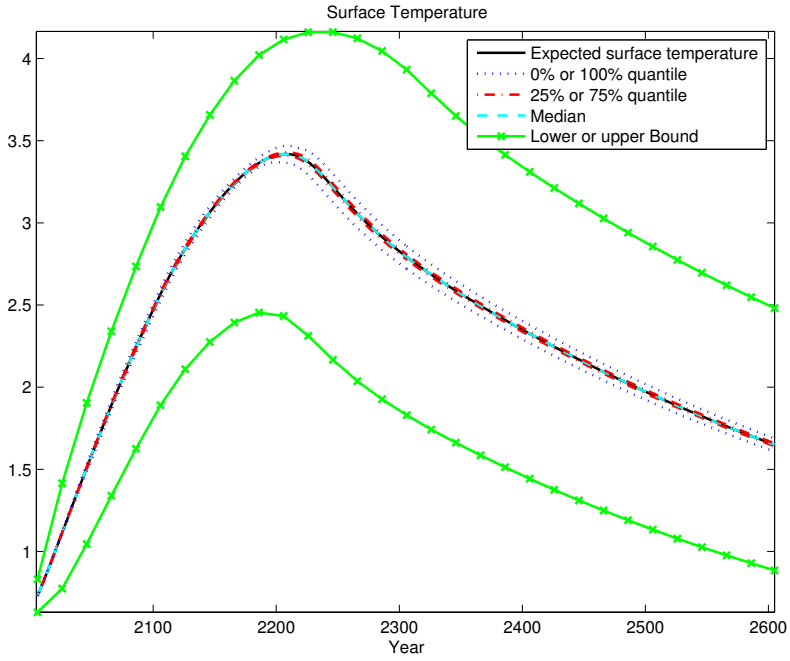

Figure 6: Carbon Tax in DSICE with an Economic Shock

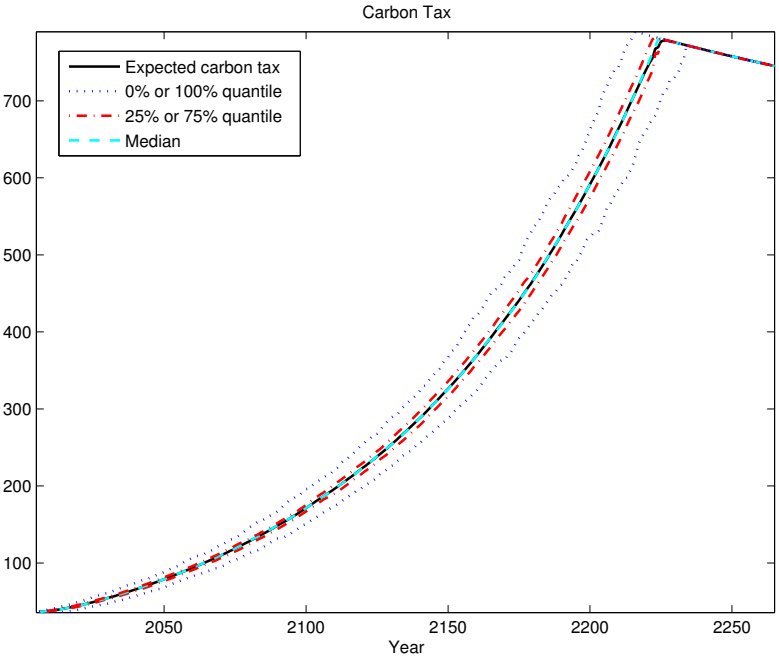

Figure 7: Capital in DSICE with a Tipping Point

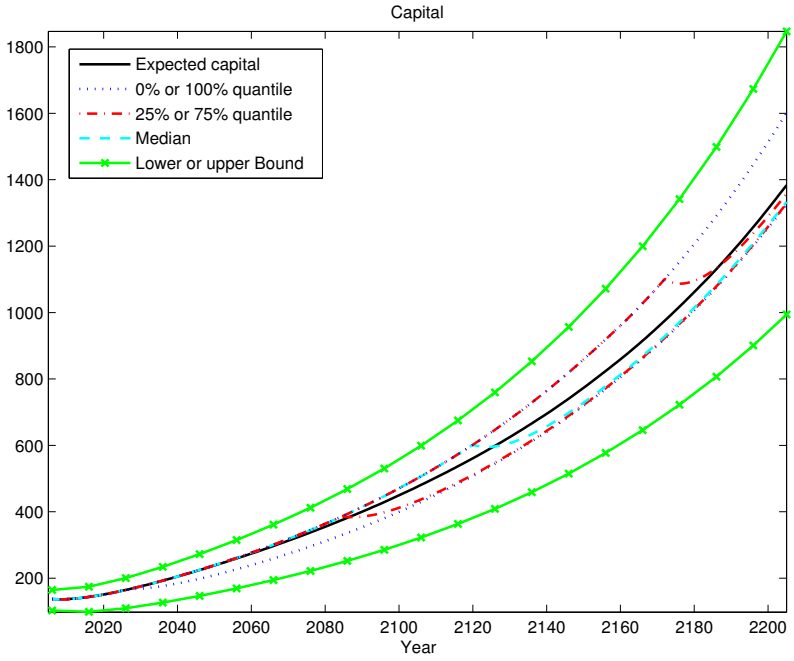

Figure 8: Atmospheric Carbon in DSICE with a Tipping Point

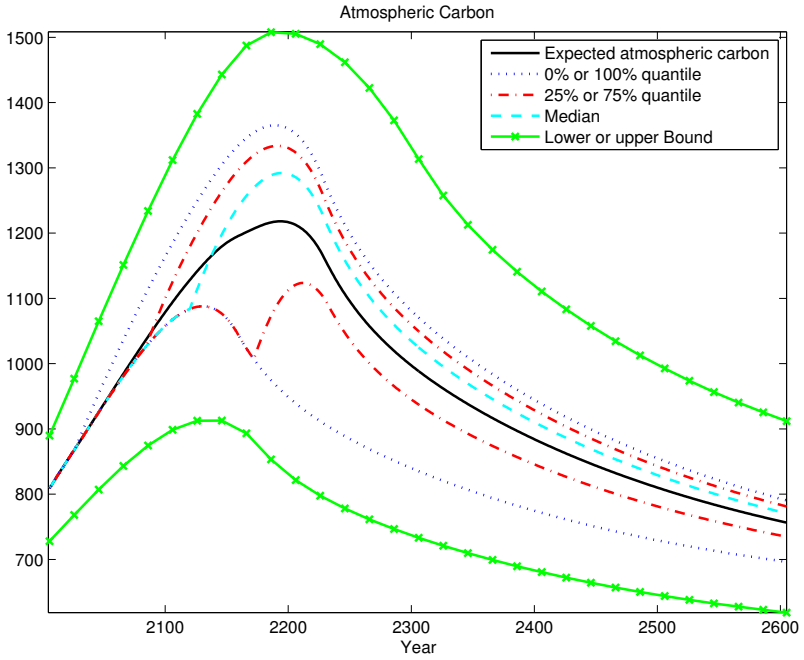

Figure 9: Surface Temperature in DSICE with a Tipping Point

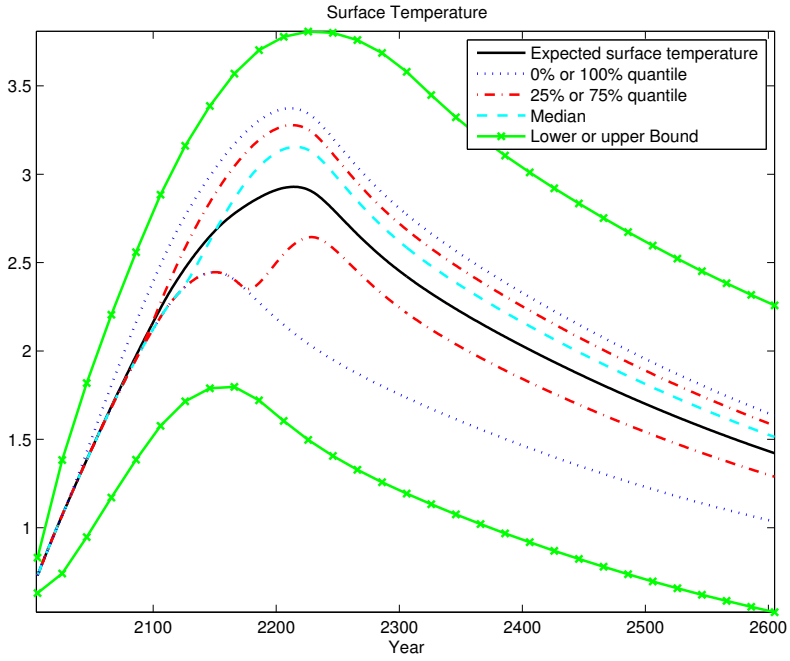

Figure 10: Carbon Tax in DSICE with a Tipping Point

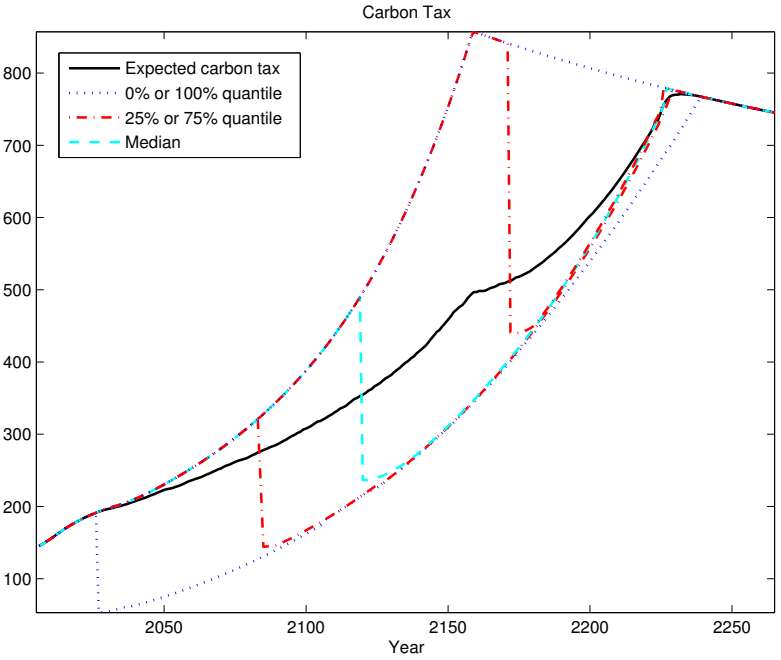

Figure 11: Capital in DSICE with an Economic Shock and a Tipping Point

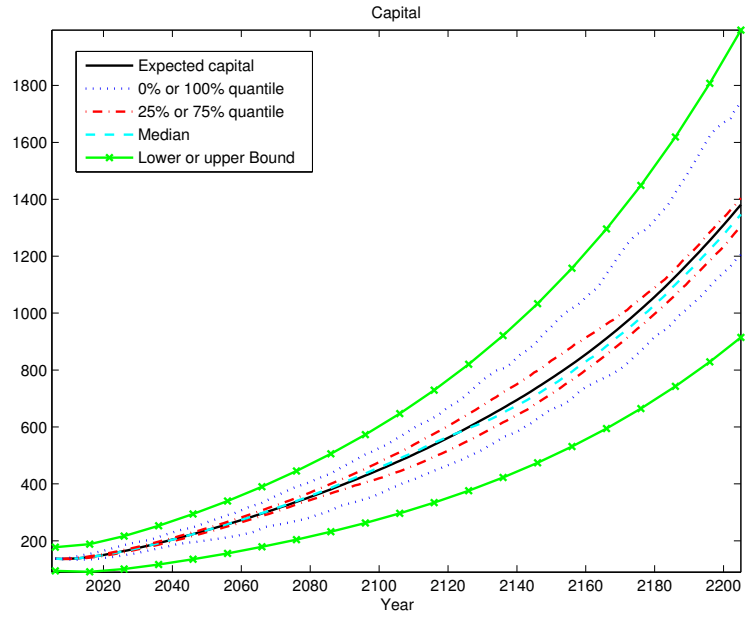

Figure 12: Carbon Tax in DSICE with an Economic Shock and a Tipping Point

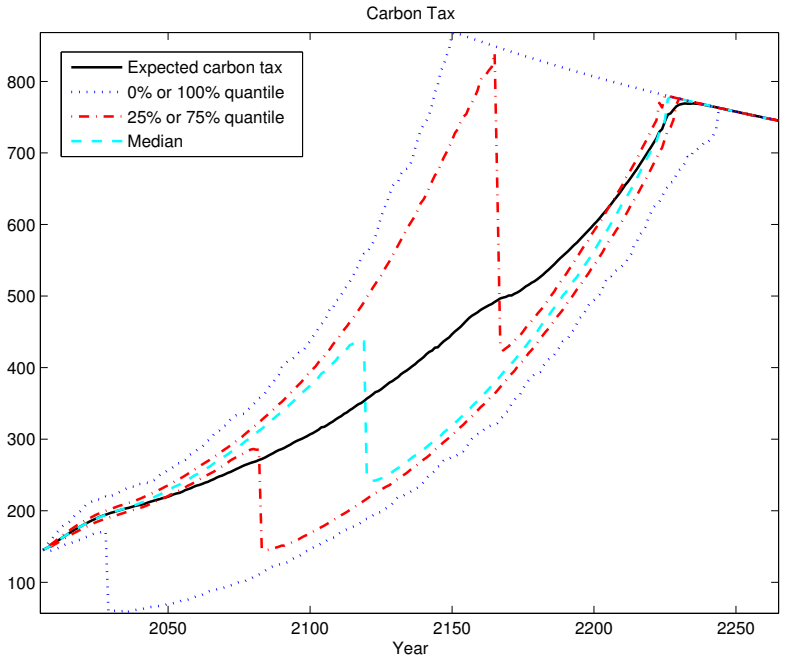

Figure 13: Carbon Tax of DSICE with Different Time Steps