University of California, Berkeley
Department of Agricultural &
Resource Economics

*CUDARE Working Papers*

*Year 2012* *Paper 1130*

# A 4-stated DICE:
## quantitatively addressing uncertainty effects in climate change

Christian Traeger

# A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change

Christian Traeger

Department of Agricultural & Resource Economics, UC Berkeley

CUDARE Working Paper 1130

**Abstract:** We introduce a version of the DICE-2007 model designed for uncertainty analysis. DICE is a wide-spread deterministic integrated assessment model of climate change. However, climate change, long-term economic development, and their interactions are highly uncertain. A thorough empirical analysis of the effects of uncertainty requires a recursive dynamic programming implementation of integrated assessment models. Such implementations are subject to the curse of dimensionality. Every increase in the dimension of the state space is paid for by a combination of (exponentially) increasing processor time, lower quality of the value function and control rules approximations, and reductions of the uncertainty domain. The paper promotes a four stated recursive dynamic programming implementation of the DICE model. Our implementation solves the infinite planning horizon problem for an arbitrary time step. Moreover, we present a closed form continuous time approximation to the exogenous (discretely and inductively defined) processes in DICE and present a Bellman equation for DICE that disentangles risk attitude from the propensity to smooth consumption over time.

**JEL Codes:** Q54, Q00, D90, C63

**Keywords:** climate change, uncertainty, integrated assessment, DICE, dynamic programming, risk aversion, intertemporal substitution, recursive utility

**Correspondence:**
Christian Traeger
Department of Agricultural & Resource Economics
207 Giannini Hall #3310
University of California
Berkeley, CA 94720-3310
E-mail: traeger@berkeley.edu

# 1 Introduction

To evaluate optimal climate policy, we have to integrate the endogenous evolution of climate into economic growth models. These so-called integrated assessment models (IAMs) are generally too complex to permit a proper incorporation of uncertainty. This limitation weighs strongly because optimal policy today depends on how a model forecasts the climate, the global economy, and their interactions over at least a couple of centuries. The present model makes a comprehensive assessment of the effects of uncertainty in the integrated assessment of climate change more accessible to a broader audience of environmental and resource economists. In contrast to Monte-Carlo approaches, our model solves for the optimal policy under uncertainty. In contrast to papers analyzing a small number of discrete probabilistic events, our model incorporates uncertainty in every period and the decision maker optimally reacts to the anticipated future resolution of uncertainty.

The DICE model by Nordhaus (2008) combines a Ramsey-Cass-Koopmans growth economy with a simple model of the carbon cycle and the climate. DICE is the most wide-spread integrated assessment model, likely so for two reasons. First, the model balances parsimony with realism. The modeler can generate realistic quantitative estimates of the optimal carbon tax without sacrificing an analytic understanding of the mechanisms driving the results. Second, DICE is an open access model that solves on an EXCEL spread sheet. We use DICE as a point of departure because its parsimony makes the effects of uncertainty more accessible to analytic introspection, and because it has the largest audience of modelers familiar with the basic model. Its parsimony implies a state space that, for replicating the full model, can be reduced to six state variables (plus time). Given the curse of dimensionality in dynamic programming, an implementation with six (or seven) state variables comes at the cost of very long processor time and significant sacrifices in the approximation to the true solution. Moreover, most economic questions introduce further state variables to capture persistence in shocks or learning. Using analytic approximations in the climate module, we cut the dimension of the state space to three, plus a time state. This reduction cuts run-time of the model by a factor of 10-100, even if we increase the number of approximating basis functions on the reduced state space tenfold.[1]

Our model differs from the original DICE-2007 model in five ways. First, we replace the difference equations describing exogenous parameters by their approximate continuous time solutions. This change makes the exogenous drivers in DICE even more accessible to introspection, and enables us to run DICE in an arbitrary

[1] The order of magnitude of processor time is hours. The precise runtime depends on the efficiency in implementing the value function iteration algorithm, the evaluation of the approximating basis functions, and the optimization software employed in the maximization step. Our cited applications of the model have run-times in the order of hours, unless we have a good initial guess. In order to avoid significantly longer run-times, implementations of the full DICE model usually employ approximations to the value function of significantly lower precision, or restrict the domain excluding rich implementations of uncertainty.

time step without recalibrating the exogenous processes. In particular, we suggest running the model in a yearly rather than decadal time step when analyzing economic fluctuations and shocks, and when modeling learning over the climate system. Second, a simplification of different heat capacity and feedback related delays approximates the evolution of atmospheric temperature in a single delay equation. Third, we approximates the carbon cycle by means of a time varying decay rate of atmospheric carbon.[2] These two approximations are crucial in reducing the state space of the model. Fourth, we normalize the model to reduce the node density required to achieve a given precision in the approximation. For this purpose, we express consumption, capital, the value function, and the Bellman equation in effective labor units. The normalization also helps the numerical model to convergence in applications to growth uncertainty. Fifth, we incorporate Epstein-Zin preferences into the welfare evaluation of uncertain scenarios. As is well-known from the finance literature, the discounted expected utility model overestimates the risk-free discount rate (risk-free rate puzzle) and underestimates risk premia (equity premium puzzle). Tuning the discounted expected utility model to get either of the discount rate or the risk premium right increases the error on the other one. Epstein-Zin preferences disentangle Arrow-Pratt risk aversion from the propensity to smooth consumption over time. They are fully rational and result in a better calibration of DICE to observed market data. In Crost & Traeger (2010) we demonstrate the relevance of Epstein-Zin preferences when evaluating climate policies with DICE.

Kelly & Kolstad (1999, 2001) and Leach (2007) implement the full DICE model as a recursive dynamic programming model, analyzing uncertainty and learning over the sensitivity of temperatures with respect to carbon emissions. These papers are impressive, careful, and efficient implementations of older versions of DICE. Apart from running at a smaller than decadal time step, our reduction of the state space increases both, speed of convergence and precision of the derived optimal policies. In particular, our accompanying applications have the necessary precision to analyze the differences that future stochasticity, or the anticipation of learning, have on today's optimal policy. These questions were likely not addressed in the earlier papers, because they require a significantly higher precision in the solution than the analysis of different learning trajectories and the statistical analysis of uncertainty resolution. A different set of papers introduces uncertainty into non-recursive implementations of integrated assessment models. Closest to our implementation, Keller, Bolker & Bradford (2004) introduce uncertainty and learning into an earlier version of DICE. Even with their highly efficient, parallized implementation on a cluster, the employed non-recursive methodology only allows for a few discrete uncertain events, or exogenous learning over three discrete state of the world realizations at one given time. For many applications, such individual uncertain events deliver interesting insights. However, these studies cannot replace comprehensive uncertainty evaluations using state

[2]We also suggest a more elaborate time and state dependent approximation that we implement in an application to tipping point uncertainty (Lemoine & Traeger 2011).

of the are stochastic dynamic programming methods. Finally, Monte-Carlo methods are the most common approach to addressing uncertainty in the integrated assessment literature. However, Monte-Carlo methods, as implemented in this strand of literature, do not model decision making under uncertainty. They present a sensitivity analysis that averages over deterministic simulations. In particular, these models cannot derive optimal policies (Crost & Traeger 2010).

Section 2 presents the continuous time approximation to the exogenous processes in DICE, the endogenous equations of motion for a freely chosen time step, and the approximations to warming delay and the carbon cycle. Section 3 derives the normalized Bellman equation, the solution algorithm, and an extension to Epstein-Zin preferences. Following a brief conclusion, the appendix discusses the calibration of the model.

# 2 The equations of motion

First, we introduce the exogenous processes in DICE. Our (approximate) continuous time solution to the iteratively defined processes in DICE make them even more amenable to introspection. Second, we introduce the heart of our DICE implementation, the endogenous equations of motions. These equations define the state transitions and take a user defined time step. Third, we explain the approximations that enable us to reduce the state space. We number (only) those equations needed for the numerical implementation.

## 2.1 Exogenous processes

Six exogenous processes derive straight from DICE-2007. We graph the recursively generated, decadal values from the original DICE model and our continuous time approximations in Figure 1. In particular, exogenous technological progress and the exogenously falling cost of abatement are main drivers of the so-called "ramp" structure of optimal abatement in DICE, i.e., the finding that abatement effort starts moderately and slowly increases over time. Two additional exogenous processes relate to our analytic approximation to the carbon cycle and the warming delay. We abbreviate growth rates by $g$ and rates of decay by $\delta$.[3]

The exogenous processes in the **economy** determine population growth, technological progress, the carbon intensity of production, and an abatement cost coefficient. Population $L_t$ simultaneously represents labor. We denote the annual growth rate of labor in period $t$ by $g_{L,t}$. The difference equations defining annual population growth

[3]A subindex indicates the growing or decaying variable. We use $\delta_A$ instead of $\delta_{g_A}$ to denote the rate of decrease of the growth rate $g_A$ of the technology level. A "*" marks a "rate" that parametrizes a speed of convergence from an initial to a final growth rate.

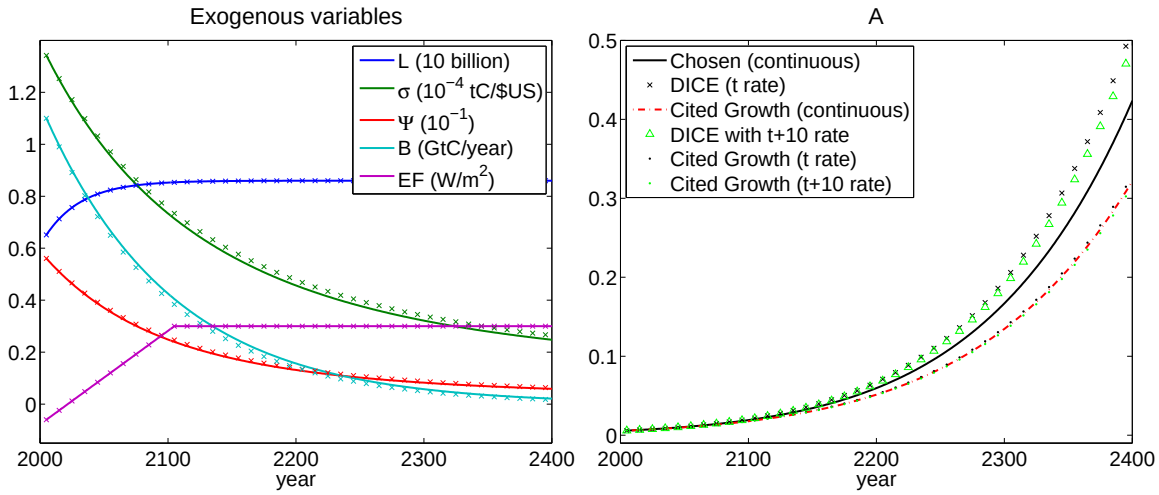


Figure 1 shows the time paths of the DICE model's exogenous parameters. The graph on the left shows our continuous time interpolations (solid) and the original, recursively defined values in the DICE model (crossed) for population (L), carbon intensity of production ($\sigma$), abatement cost coefficient ($\Psi$), emissions from land-use change and forestry (B), and exogenous forcing (EF). The graph on the right shows the labor productivity (A) in DICE (crossed), the cited labor productivity (dash-dotted), and our chosen interpolation (solid). The "t+10 rate" curves grow productivity with the growth rate prevailing at the end of the according decade, while "t rate" curves grow labor productivity with the growth rate prevailing at the end of the according decade. Together, they bound a corresponding continuous model.

in DICE have the continuous time approximation[4]

$$g_{L,t} = \frac{g_L^*}{\frac{L_\infty}{L_\infty - L_0}\exp[g_L^*\ t] - 1}\ , \tag{1'}$$

corresponding to the analytic continuous time solution characterizing period $t$ population

$$L_t = L_0 + (L_\infty - L_0)(1 - \exp[-g_L^*\ t])\ . \tag{2}$$

Here, $L_0$ denotes the initial and $L_\infty$ the asymptotic population. The parameter $g_L^*$ characterizes the speed of convergence from initial to asymptotic population.

The technology level $A_t$ in the economy grows at an exponentially declining rate

$$g_{A,t} = g_{A,0}\exp[-\delta_A t]\ , \tag{3'}$$

leading to the analytic continuous time solution

$$A_t = A_0 \exp\left[g_{A,0}\frac{1 - \exp[-\delta_A\ t]}{\delta_A}\right]\ . \tag{4}$$

[4]We mark two equations below with a prime (1′ and 3′). These are not needed to evaluate the exogenous processes; they can be used to evaluate the discount factor $\beta_t$ in the renormalized Bellman equation as we explain in section 3.

The parameter $g_{A,0}$ denotes the initial growth rate and $\delta_A$ its constant rate of decline. While the original DICE model employs a Hicks-neutral formulation of technological change, we use labor augmenting technological progress. The two are equivalent under the given Cobb-Douglas production. However, the more wide-spread labor augmenting formulation simplifies the reformulation of the model in effective labor units, which increases the implied node density for capital. This reformulation improves the numerical approximation for a given number of nodes in the capital dimension.[5]

The original DICE model calculates the technology level as $A^{DICE}_{t+10} = \frac{A^{DICE}_t}{1-10g_{A,t}}$, using 10 year time steps and the approximation $\frac{1}{1-g} \approx 1+g$. However, the approximation $\frac{1}{1-10g_{A,t}}$ is significantly larger than $1+10g_{A,t}$ or $\exp(10g_{A_t})$. Therefore the effective growth rate employed in the DICE model is significantly larger than $g_{A,t}$. The dash-dotted line in Figure 1 depicts the continuous time productivity path that actually replicates the initial decadal growth rate $g^{DICE}_{A,0}$ stated as the DICE reference. The crossed decadal values correspond to the technology levels created in DICE because of the growth rate approximation. We decided to use $g^{DICE}_{A,0}/10$ as the annual growth rate in our continuous time model, which generates the solid line between the original DICE path (crossed) and the growth path matching the referenced value $g^{DICE}_{A,0}$ (dash-dotted). The additional lines discuss consequences of the decadal time step.[6]

The DICE model assumes an exogenous decrease of the carbon intensity of production. The decarbonization factor of production grows at the (decreasing) rate $g_{\sigma,t} = g_{\sigma,0}\exp[-\delta_\sigma\ t]$, leading to the continuous time representation

$$\sigma_t = \sigma_0 \exp\left[g_{\sigma,0}\frac{1-\exp[-\delta_\sigma\ t]}{\delta_\sigma}\right]\ . \tag{5}$$

In addition, the economy can pay for abating emissions. The abatement cost coeffi-

[5]This reformulation also increases the numerical stability when modeling growth uncertainty. Finally, labor augmenting technological progress implies a balanced growth for a broader class of production functions in the Ramsey-Cass-Koopmans model.

[6]The growth rate in DICE falls over time. In the original DICE model, the growth rate at the beginning of a decade generates growth throughout the decade, which generates more technological progress then with a smaller (or continuous) time step. The triangles bound this second, technological progress increasing effect by showing DICE's evolution of the technology level if we used the growth rate at the end of a given decade to generate growth. Similarly, the dots just above and just below the dash-dotted calibration line show the effects of using a decadal time step. The dots just above (below) the line use the growth rate at the beginning (end) of a decade, instead of a continuous model. Together, the additional lines point out that this discretization effect is very minor with respect to the approximation effect of the growth rate, which generates the difference between the calibration line and the original DICE line.
Note that we depict the technology level in terms of labor augmenting technological progress. The equivalent growth rates in terms of total factor productivity, employed in the original DICE model, are lower by the factor $1-\kappa$. Hence, the rounding error is slightly lower when using total factor productivity growth rates, an effect that our triangled curve (and the original DICE curve) take into account.

cient $\Psi_t$ falls exogenously over time and is given by

$$\Psi_t = \frac{\sigma_t}{a_2} a_0 \left( 1 - \frac{(1 - \exp[g^*_\Psi \, t])}{a_1} \right) . \tag{6}$$

The parameter $a_0$ denotes the initial cost of the backstop (in 2005), $a_1$ denotes the ratio of initial over final backstop,[7] and $a_2$ denotes the cost exponent (see also equation 13 below). The rate $g^*_\Psi$ captures the speed of convergence from the initial to the final cost of the backstop.

The exogenous processes on the **climate** side of DICE govern non-industrial $CO_2$ emissions and radiative forcing[8] from non-$CO_2$ greenhouse gases. In addition, our state space reduction introduces an exogenous process governing the removal of excess carbon from the atmosphere and the cooling due to the ocean's heat capacity. DICE assumes an exponential decline of $CO_2$ emissions from land use change an forestry

$$B_t = B_0 \exp[-\delta_B \, t] . \tag{7}$$

Non-$CO_2$ greenhouse gases are exogenous to the model and cause the radiative forcing

$$EF_t = EF_0 + 0.01(EF_{100} - EF_0) \times \min\{t, 100\} . \tag{8}$$

Note that exogenous forcing starts out slightly negatively.

Our approximation of DICE's carbon cycle uses an exogenous removal rate $\delta_{M,t}$ of atmospheric $CO_2$ that is in excess of preindustrial levels

$$\delta_{M,t} = \delta_{M,\infty} + (\delta_{M,0} - \delta_{M,\infty}) \exp[-\delta^*_M \, t] . \tag{10}$$

Moreover, we use an exogenous estimate of the atmosphere-ocean temperature differential

$$\Delta T_t = \max\{0.7 + 0.02t - 0.00007t^2, 0\} , \tag{12}$$

which governs transient cooling of the atmosphere caused by the oceans' heat capacity.

## 2.2 Endogenous equations of motion

Apart from time, the model features three state variables: produced capital $K_t$, the stock of atmospheric carbon $M_t$, and temperature $T_t$. Temperature is a state variable because atmospheric warming happens with a delay: the heat capacity of the ocean and various feedback processes delay the temperature increase. Therefore, next

[7] The general interpretation is more precisely that $a_1$ is the ratio $\frac{initial\ cost\ of\ backstop}{initial\ cost\ of\ backstop - final\ cost\ of\ backstop}$. However, for the employed value of 2 both ratios are the same, so we stick with Nordhaus's interpretation.

[8] Radiative forcing is a measure for the change in the atmospheric energy balance. The reader may think of it as the flame that greenhouse gases turn on to slowly warm the planet over time.

period's temperature depends not only on the atmospheric carbon concentrations, but also on current temperature. We follow Kelly & Kolstad (1999) in incorporating time as a state variable, which makes it possible to contract the Bellman equation to an arbitrary precision despite the intrinsic non-stationarity of the DICE model. Moreover, the time state enables us to solve the model for an infinite time horizon with an arbitrary time step. Here, we present the deterministic equations of motions replicating DICE-2007. Applications, cited in the introduction, introduce uncertainty into the equations of motion. In addition, persistent shocks or learning will introduce additional equations of motion governing informational states.

Capital $K_t$, labor $L_t$, and labor augmenting technology $A_t$ enter the Cobb-Douglas production function, turning into gross (or potential) output $Y_t^{gross} = (A_t L_t)^{1-\kappa} K_t^{\kappa}$. The parameter $\kappa$ is the income share of capital. Driven by technological progress and population growth, capital grows by an order of magnitude over the next century. Approximating $K_t$ on a time constant grid would either require an excessive amount of nodes, or imply a crude approximation. We therefore follow the macroeconomic tradition expressing consumption and capital in per effective labor units $c_t = \frac{C_t}{A_t L_t}$ and $k_t = \frac{K_t}{A_t L_t}$. Then, gross production in effective labor units is $y_t^{gross} = \frac{Y_t^{gross}}{A_t L_t} = k_t^{\kappa}$. We introduce a flexible time step $\Delta t$, keeping our flow variables, including consumption and production, defined in units per year. Then, production during the period $[t, t + \Delta t]$ is $k_t^{\kappa} \Delta t$. The model calibration in the appendix uses $\Delta t = 1$.

The transformation of gross output into net output defines the interface between the Ramsey-Cass-Koopmans economy and the climate system. Net production follows from gross production by subtracting abatement expenditure and climate damages

$$y_t = \frac{1 - \Lambda(\mu_t)}{1 + D(T_t)} \; k_t^{\kappa} = \frac{1 - \Psi_t \mu_t^{a_2}}{\left(1 + b_1 T_t^{b_2}\right)} \; k_t^{\kappa} \; . \tag{13}$$

The function

$$\Lambda(\mu_t) = \Psi_t \mu_t^{a_2}$$

characterizes abatement expenditure as a fraction of gross output. It is a function of the emission control rate $\mu_t \in [0, 1]$ (abatement rate). This abatement rate characterizes the percentage of emissions avoided under a climate policy, as compared to a laissez-faire world.

As a fraction of gross output, the damage function

$$D(T_t) = b_1 T_t^{b_2}$$

reduces net production as a consequence of the temperature increase $T_t$ over temperatures in 1900. Net production not consumed is invested in capital, implying the equation of motion

$$k_{t+\Delta t} = \left[\left(1 - \delta_k \Delta t\right) k_t + y_t \Delta t - c_t \Delta t\right] \exp[-(g_{A,t} + g_{L,t})\Delta t] \; , \tag{15}$$

where $\delta_K$ is the annual rate of capital depreciation. The exponential function is a consequence of expressing capital in effective labor units; it reflects that the normalizing effective labor units grow by $g_{A,t} + g_{L,t}$ from one year to the next.

Anthropogenic emissions are the sum of industrial emissions and emissions from land use change an forestry $B_t$

$$E_t = (1 - \mu_t)\, \sigma_t A_t L_t k_t^{\kappa} + B_t \ . \tag{16}$$

Industrial emissions are proportional to gross production $A_t L_t k_t^{\kappa}$, and the emission intensity of production $\sigma$, and they are reduced by the emission control rate $\mu_t$. The flow of $CO_2$ emissions accumulates in the atmosphere. Atmospheric carbon in the next period is the sum of preindustrial carbon $M_{pre}$, current excess carbon in the atmosphere $M_t - M_{pre}$ net of its (natural) removal, and anthropogenic $CO_2$ emissions

$$M_{t+\Delta t} = M_{pre} + (M_t - M_{pre})\,(1 - \delta_{M,t}\Delta t) + E_t \Delta t \ . \tag{17}$$

The pre-industrial emission stock $M_{pre}$ is the steady state level in the absence of anthropogenic emissions. Equation (17) is our approximation to the carbon cycle in DICE-2007.

The atmospheric temperature change is a delayed response to radiative forcing

$$F_{t+\Delta t} = \eta_{forc} \frac{\ln \frac{M_{t+\Delta t}}{M_{preind}}}{\ln 2} + EF_t \ ,$$

which is the sum of the forcing caused by atmospheric $CO_2$ and the non-$CO_2$ forcing that follows the exogenous process $EF_t$. Note that the forcing parameter $\eta_{forc}$ contains the climate sensitivity parameter, which characterizes the equilibrium warming response to a doubling of preindustrial $CO_2$ concentrations. The temperature state's equation of motion is

$$T_{t+\Delta t} = (1 - \sigma_{forc}) T_t \ \Delta t + \sigma_{forc} \ \frac{F_{t+\Delta t}}{\lambda} \Delta t - \sigma_{ocean} \ \Delta T_t \ \Delta t \ . \tag{18}$$

The last term replaces the oceanic temperature state in DICE-2007.

The appendix calibrates the model to reproduce the original policies of DICE-2007 with a one year time step. Running the model with time steps between a few months and a few years will result in only minor differences in the optimal policy. A word of caution about scaling the time step to arbitrary levels without re-calibrating the model is in place. Reducing the time step in the endogenous equations of motions implies what best can be summarized as a "compound interest" effect. Cutting the time step into half, cuts production (and carbon decay) per period into half. Production grows over time, and capital enters the production function. Therefore, the economy will have slightly more capital available at the end of the two "half periods" than under a full time step. Similarly, carbon decay will not be exactly the same.[9] Our

[9]Our exogenous rate of removal of excess carbon seems to be a bit more robust than scaling the time step in the actual carbon cycle.

formulation of the model fixes the evolution of all exogenous processes to a continuous time approximation to DICE-2007. Thus, as opposed to the original formulation of these processes as difference equations, we avoid that changes in the time step change the exogenous input of our model. The time step dependence of the endogenous equations of motions, however, cannot be avoided. DICE-2007 itself is calibrated to a decadal time steps. Thus, when downscaling DICE-2007 from its decadal step to an annual or even monthly time step, we have to recalibrate the model if we want to match the same calibration data.[10]

## 2.3 Discussion of the state reducing approximations

Our model replaces the **carbon cycle** by a time-dependent rate of carbon removal from the atmosphere. Appendix B calibrates this rate of carbon removal to the full (deterministic) DICE model. Our approximation matches the deterministic baseline well, and we discuss how it performs when the optimal response to uncertainty implies a deviation from the deterministic policies.

Our formulation in terms of a rate of a removal of excess carbon is particularly stable across different carbon trajectories. Atmospheric carbon does not decay chemically, but it eventually moves to reservoirs other than the atmosphere. Along a path of increased emissions, reservoirs would fill up slightly faster. However, partially offsetting the consequences of the faster "saturation" is that, at the same time, a higher emission trajectory implies a higher partial pressure of carbon in the atmosphere, pushing more carbon into the reservoirs. Simulations with the full DICE model show that even largely different abatement scenarios imply almost the same effective decay rate of excess carbon for at least the first century. The effective decay rates only start to differ more notably when we approach peak carbon concentrations, under optimal policy one to two centuries in the future.

We argue that the policy impact of approximating the carbon cycle is of third order. Uncertainty causes a first order deviation of the policy. This first order deviation slowly changes the stock of carbon in the atmosphere and in other reservoirs. The evolving change in the *difference* of carbon concentrations between atmosphere and other carbon reservoirs causes a change of the rate of carbon removal, which is a second order effect. This second order effect is what we observed to be very small for the first century, even in the case of very large first order policy changes. The second order change in the decay rate has a flow effect and, only over time, grows into a notable stock (concentration) effect that will impact warming. The relation be-

[10]In addition to these "compound interest" effects, the finer time step permits a more finely tuned policy. This additional freedom in setting policy will lower the social cost of carbon slightly. We discuss the effect in the appendix to Lemoine & Traeger (2011). Quantitatively, this fine tuning effect on policy is smaller than the "compound interest" effect. In order to make our policies directly comparable to DICE-2007, we decided to calibrate our optimal policies directly to DICE-2007. See Lemoine & Traeger (2011) how the modeler can alternatively fix the policy step to 10 years, when calibrating the model for an arbitrary time step.

tween carbon concentration and temperature is again governed by a delay equation, delaying warming and economic impact once more in the order of decades. Thus, our approximation of the carbon cycle causes a rather moderate approximation error in far future impacts. Under the usual economic discounting regimes, these have a very minor impact on welfare and, thus, optimal policy. We argue that this small approximation error is usually much more than offset by gains in numerical precision when enlarging the basis for the value function approximation. Our carbon cycle approximation drops two state variables. For each dropped state we can approximately triple the number of basis functions in the remaining state space and still cut processor time by half. We will generally obtain a much more accurate solution in less time. Moreover, in difference to the numerical approximation error when using a small set of basis functions, our approximation is analytic and more easily amenable to introspection.

In addition, our model replaces a set of **temperature delay equations** in DICE by a single approximate delay equation. This step drops ocean temperature as a state variable, and approximates ocean cooling by a time dependent cooling process. We can merge the system of delay equations in DICE-2007 into the two equations

$$\begin{aligned} T_{t+10} &= T_t + C1\left[F_{t+10} - \lambda T_t + C3(T_t^{Ocean} - T_t)\right] \\ T_{t+10}^{Ocean} &= T_t^{Ocean} + C4(T_t - T_t^{Ocean}) = (1 - C4)T_t^{Ocean} + C4\,T_t \ , \end{aligned}$$

which govern atmospheric and oceanic temperatures. We find that an exogenous estimate of the difference between oceanic temperatures and atmospheric temperature gives a reasonable approximation for the evolution of atmospheric temperatures and a very good approximation of the evolution of optimal policies. For this purpose, we rewrite the first equation as

$$T_{t+10} = (1-\sigma_{forc}^{dec})T_t + \sigma_{forc}^{dec}\ \frac{F_{t+10}}{\lambda} - \sigma_{ocean}^{dec}\Delta T_t \ ,$$

where $\Delta T_t = T_t - T_t^{Ocean}$ is the ocean-atmosphere temperature difference. The parameters $\sigma_{forc}^{dec} = C1 * \lambda$ and $\sigma_{ocean}^{dec} = C1 * C3$ are decadal lag parameters, governing how atmospheric temperatures adjust to radiative forcing and to oceanic temperature. To obtain equation (18), we downscale the equation to a one year time step. A one to one mapping to our finer time resolution is not possible because the model is non-stationary. Thus, we calibrate the parameters $\sigma_{forc}$ and $\sigma_{ocean}$ in equation (18) to obtain the best fit to the deterministic DICE model (Figure 4).[11]

[11] Assuming constant forcing and feedbacks, we find an equation of the form $T_{t+10} = (1-\sigma_{forc}^{dec})T_t + \Gamma$ (for some $\Gamma \in \mathbb{R}$), and downscaling of the decadal delay parameters to the one year time step would result in $\sigma_{forc} = 1-(1-\sigma_{forc}^{dec})^{\frac{1}{10}} \approx 0.032$. For $\sigma_{forc}$, this theoretically derived approximation indeed returns the best calibration.

Note as well that, by merit of the equation governing oceanic temperatures, our exogenous approximation of the ocean-atmosphere temperature gradient is equivalent to an exogenous approximation of the temperature difference $T_t^{Ocean} - T_{t+10}^{Ocean}$ in the oceans between two subsequent periods.

Instead of the endogenous equation for oceanic temperature, we use the exogenous process for $\Delta T_t$ stated in equation (12) and depicted in Figure 5. The simple quadratic-max interpolation in equation (12) captures $\Delta T_t$ in the baseline run almost perfectly for the first 150 years. Later, it deviates notably until both the true temperature difference and its approximation converge to zero in equilibrium. Our model calibrations in the appendix (e.g. Figure 2) reflect a small deviation of our atmospheric temperature path from the one in the original DICE model during the second 150 years. This deviation is an immediate consequence of this simple to implement, but for the distant future slightly crude approximation. The impact on optimal policies is minimal. We briefly comment on a more sophisticated approximation below. The relevant question is how the approximation reacts to policy changes caused by introducing uncertainty. First order policy changes imply a change in the emission flow that builds up into a change of the emission stock, which changes radiative forcing. The change in radiative forcing will warm both atmosphere and ocean, but the atmosphere slightly more quickly, temporarily resulting in a second order difference in $\Delta T$. Even when doubling the abatement effort, simulations show that $\Delta T$ only starts to differ notably more than half a century in the future. Then, this difference starts affecting atmospheric temperature, proportional to the delay parameter $\sigma^{dec}_{ocean} \approx 0.007$, which keeps the effects on atmospheric temperatures small and delayed. Given the usual discounting, the policy feedback of these differences will be second if not third order.

The modeler may want to selectively employ only the carbon cycle approximation, or only the temperature delay equation approximation. When eliminating approximations, our recommendation is to adopt the carbon cycle approximation, but implement ocean temperature endogenously. The latter approximation is slightly more crude and less easy to evaluate. However, we think that the full implementation of ocean temperature evolution generally does not warrant an additional state at the expense of reducing the order of the approximating basis functions. A way of improving the warming delay model without reintroducing a state variable replaces the simple quadratic approximation in equation (12) by a more complex numerical interpolation of the actual DICE output for the closest deterministic scenario.[12]

Finally, note that we can further improve the approximations of the carbon cycle and the warming delay by employing multivariate interpolations. Using e.g. a multidimensional Chebychev basis, we can interpolate the decay rate and the temperature difference between atmosphere and ocean as a function not only of time, but also the other state variables. The current carbon stock and the current temperature

[12]We tried a large set of simple closed form approximations, and for the first 150 years our linear-quadratic approximation combined with the minimum clearly performed best. However, a simple spline or higher order Chebychev interpolation will obviously produces better overall results. Implementing such a higher order interpolation, we recommend that the modeler pays attention to the interpolation converging to zero, e.g. by using exponential damping. Our intention here is to write down an easy to implement model and we leave a fitting to arbitrary accuracy of the exogenous process with basis functions of his choice to the reader.

carry information on the historic emission path and, thus, the changes induced for the effective carbon removal rate and ocean cooling when deviating from the baseline. For that purpose, we generate a set of deterministic DICE-2007 paths with different emission policies in the neighborhood of the baseline. Then, we interpolate not only the baseline policy, but also the effects induced by deviating from the deterministically optimal policy. See Lemoine & Traeger (2011) for an application.

# 3 Welfare and Bellman equation

We first explain the dynamic programming problem using standard preferences. Then, we discuss the solution algorithm. Finally, we present the comprehensive Bellman equation for Epstein-Zin-Weil preferences that disentangle risk aversion from the propensity to smooth consumption over time.

## 3.1 Bellman equation for standard preferences

An optimal decision under uncertainty has to anticipate all possible future realizations of the random variables together with the corresponding optimal future responses. The Bellman equation reduces the complexity of the decision tree by breaking it up into a trade-off between current consumption utility and future welfare, where future welfare is a function of the climatic and economic states in the next period. This so-called value function $V(K_t, M_t, T_t, t)$ characterizes the maximal expected welfare a decision maker can derive over the infinite time horizon, given a particular state of the economy. The optimization problem is essentially solved once we find an approximation to $V$.

The welfare flow in a period is the population $L_t$ weighted utility from per capita consumption $\frac{C_t}{L_t}$. DICE uses a constant intertemporal elasticity of substitution, and we denote its inverse, the coefficient of aversion to intertemporal change, by $\eta$. We denote the rate of pure time preference, also called utility discount rate, by $\delta_u$. The Bellman equation for standard (entangled) preferences is

$$V(K_t, M_t, T_t, t) = \max_{C_t, \mu_t} L_t \frac{\left(C_t/L_t\right)^{1-\eta}}{1-\eta} \Delta t \; + \\ \exp[-\delta_u \Delta t] \; \mathrm{E} \; V(K_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, t+\Delta t),$$

where E takes expectations over uncertainty in the equations of motion governing period $t + \Delta t$ states (introduced in the stochastic applications of the model). The right-hand side of the Bellman equation describes the optimization problem in period $t$. The optimal decision maximizes the sum of immediate consumption utility and the discounted expected value of future welfare. This maximization is subject to the equations of motions (15), (17), and (18), or their modifications including

uncertainty, and the constraints

$$0 \leq \mu_t \leq 1 \text{ and } 0 \leq C_t \leq Y_t \ .$$

As we pointed out in section 2.2, approximating the value function over $K_t$ would be computationally inefficient, because capital levels change significantly over time. Using effective labor units for measuring consumption, capital, and production, we also define a new value function measuring the value of the optimal program in units closely resembling effective labor

$$V^*(k_t, M_t, T_t, t) = \left. \frac{V(K_t, M_t, T_t, t)}{A_t^{1-\eta} L_t} \right|_{K_t = k_t A_t L_t} \ .$$

Then, Appendix A transforms the Bellman equation into the dynamic programming equation

$$V^*(k_t, M_t, T_t, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} \Delta t + \frac{\beta_{t,\Delta t}}{1-\eta} \, \mathrm{E}\left[V^*(k_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, t+\Delta t)\right]. \quad (19)$$

The parameter $\beta_{t,\Delta t}$ defines the growth adjusted discount factor

$$\beta_{t,\Delta t} = \exp\left[(-\delta_u + g_{A,\tau}(1-\eta) + g_{L,\tau})\Delta t\right] \ . \quad (20)$$

Its time dependence arises because of the non-constant growth rates in DICE. The factor determines the contraction of the Bellman equation.[13] We can either use the exact, time-step dependent formulas for the growth rates

$$g_{A,\tau} = [\ln(A_{t+\Delta t}) - \ln(A_t)]/\Delta t \quad \text{and} \quad (1^*)$$

$$g_{L,\tau} = [\ln(L_{t+\Delta t}) - \ln(L_t)]/\Delta t \ , \quad (3^*)$$

or their continuous time approximations $(1')$ and $(3')$ at the expense of a small error.[14]

We derive the normalized Bellman equation (19) in Appendix A for the general case of Epstein-Zin preferences. A useful "side-effect" of solving for the value function,

[13] For too low a time preference relative to growth and intertemporal substitutability, the Bellman equation will not contract. Practical convergence problems can already arise before expected welfare diverges and makes the maximization problem theoretically ill-posed. Note that the time dependence of $\beta_{t,\Delta t}$ is entirely a consequence of the non-constant growth rates and does not imply a time inconsistent objective function.

[14] The error is bounded by the change of the growth rate. A more precise evaluation of the differences results in a truly negligible error for $g_{A,\tau}$, changing the discount factor in the order $10^{-6}$ for an annual time step. However, the initially quickly growing labor can imply an approximation error in the continuous time formula of up to a percent of the discount factor, i.e., of the order $10^{-4}$ in the discount factor. That error is still small as opposed to any knowledge we have with respect to the true discount factor, but we can avoid it using equation $(3^*)$ instead.

rather than just an optimal path, is that we obtain the social cost of carbon directly as

$$SCC_t = -\frac{\partial_{M_t} V}{\partial_{K_t} V} = -\frac{\partial_{M_t} V^*}{\partial_{k_t} V^*}\ A_t\ L_t\ .$$

With this formula, we can calculate the social cost of carbon even when full abatement is achieved, i.e., we can calculate the value of carbon sequestration from the atmosphere also after the abatement rate hits the constraint.[15]

## 3.2 Solving the model

For the numerical implementation, it is usually more efficient to maximize over the abatement cost $\Lambda$ rather than over the abatement rate $\mu$. The two are strictly monotonic transformation of each other and (only) the constraints on $\Lambda$ are linear

$$\begin{pmatrix} 1 & \dfrac{k_t^\kappa}{1+b_1 T_t^{b_2}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} c \\ \Lambda \end{pmatrix} \leq \begin{pmatrix} \dfrac{k_t^\kappa}{1+b_1 T_t^{b_2}} \\ \Psi_t \end{pmatrix} \quad \text{and } c, \Lambda \geq 0\ . \tag{21}$$

Apart from the physical states, we have to approximate the value function over state variable $t$ on the interval $[0, \infty)$. Its natural unboundedness is inconvenient when generating the approximation grid. It is helpful to introduce a strictly monotonic transformation that maps $t \in [0, \infty)$ to

$$\tau = 1 - \exp[-\zeta t] \in [0, 1)\ .$$

We refer to $\tau$ as artificial time. We then generate the grid on the time axis using e.g. Chebychev nodes on $[0, 1]$.[16] A larger choice of the numerical parameter $\zeta$ moves the nodes closer to the early (real-time) periods.[17] After generating the nodes, we transform them back to real time using the inverse transformation $t = -\frac{\ln[1-\tau]}{\zeta}$.

[15] DICE-2007 obtains the social cost of carbon indirectly from the condition that it has to equal the optimal abatement cost. Once full abatement is achieved, usually some time during the next century, such an approach can no longer track the actual social cost of carbon.

[16] Chebychev nodes do not lie on the boundary of the interval.

[17] The logarithmic transformation clusters nodes densely in early as opposed to late periods. Such a clustering is useful for the DICE model, where most of the action happens relatively early in the time horizon, when the exogenous processes exhibit the highest rates of change, and when the economy transitions from a high emission to a low emissions path. However, a simple logarithmic transformation would exaggerate such clustering. A low parameter $\zeta$ is necessary to moderate the clustering at early times and spread the support to a sufficient range in real time. The parameter $\zeta$ is a purely numeric parameter and we chose $\zeta = 0.02$ for the runs depicted in the appendix. In general, we suggest smaller values rather than increasing the parameter. Sometimes it will be worthwhile playing with the parameter spreading nodes differently in order to improve the value function fit or even convergence properties.

The challenge is to find a reasonable approximation to the true value function. The solution technique relies on Bellman's observation that, in our infinite horizon setting, the value function on the left and on the right of equation (19) coincide. We obtain a solution to equation (19) by function iteration, approximating the value function $V$ by a set of basis functions. We use and recommend Chebychev polynomials. A frequently used alternative are cubic splines. The compecon toolbox provided by Miranda & Fackler (2002) provides convenient tools for the function approximation step, requiring little to no knowledge of the underlying theory. The following steps outline the algorithm.

1) Setting up the problem:

1. Choose the intervals on which to approximate the value functions. We suggest such intervals in the appendix. In general, a reasonable interval choice depends on the type and magnitude of the modeled uncertainty.
2. Choose an approximation method for the value function $V$. Our calibration of the model uses Chebychev polynomials.
3. Generate a grid on the product space of the approximation intervals. When using Chebychev polynomials, use Chebychev nodes.[18]
4. Start with an arbitrary guess for the value function or the coefficients characterizing its approximation.

2) The function iteration:

1. Solve the right hand side optimization problem of the Bellman equation (19) for every point $i$ on the grid. Save the optimal control variables $c(i)$, $\mu(i)$ (or $\Lambda(i)$), and the maximized objective $v(i)$ for every point on the grid.
2. Fit a new value function approximation using the newly generated values $v(i)$. This new fit generates new basis coefficients $g(i)$.
3. Check whether the change in coefficients or values satisfies a given tolerance criterion. If yes, stop. If not, use the new coefficients $g(i)$ returning to step 1, employing the optimized controls from the previous iteration as initial guesses for the maximization problem.

Given the (approximate) value function, we can analyze the control rules and simulate different representations of the optimal policy over time. For the simulation, we either fit a continuous control rule, or we forward-solve the Bellman equation, knowing the value function, starting from the initial state. Under uncertainty, we can quickly simulate a large set of runs and depict statistical properties. After a first solution, we recommend checking whether changes in tolerance, number of approximating basis

[18] A simple an efficient basis for a multi-dimensional state space is the tensor basis. It contains the tensor product of all combinations of basis functions in the different dimensions. The corresponding grid and basis is automatically generated when using the compecon toolbox. Sometimes, it is suggested to drop higher order cross terms as a way of saving basis functions or nodes. The loss in approximation quality of this approach strongly depends on the precise form of the value function.

functions, and changes in the interval bounds still affect our results. Plotting of the control rules, more precisely their two or three dimensional cuts, often gives a good idea of the approximation quality. Using Chebychev polynomials, a low order of the basis generally results in major wave patterns in the control rules. These patterns are entirely of numeric origin. Various modifications of the basic function iteration algorithm described above can be used to speed up convergence. The most important "trick" to speed up the algorithm is known as Howard's method, or modified policy iteration: after every value function iteration, Howard's method iterates the Bellman equation several times without without re-optimizing the controls. We emphasize, first, that the algorithm solves the problem on the continuous state and control space. Second, by making time a state variable, our approach contracts to the true solution without depending on the initial guess.

We can reduce the state space to only 3 dimensions, if we are willing to step back discretely in time from a finite planning horizon. The solution algorithm is similar to the one described above. However, it becomes more important to start with a good initial guess, because this guess directly determines optimal policies close to the end of the planning horizon. The time horizon should therefore be at least several centuries. The modeler should test different initial guesses and compare the solutions. As long as the result stays sensitive to the initial guess, he should push the time horizon further out.[19]

## 3.3 Epstein-Zin preferences

The Bellman equation in section 3.1 reflects the discounted expected utility model. A serious short-coming of this model is its inability to correctly capture the risk-free discount rate and risk premia. Basal & Yaron (2004) show how Epstein-Zin preferences explain the corresponding equity premium and the risk-free rate puzzles in finance. Their approach builds on a model by Epstein & Zin (1989) and Weil (1990) that disentangles risk aversion from a decision maker's propensity to smooth consumption over time. Note that these preferences are fully rational (Traeger 2010), in particular they obey the von Neumann & Morgenstern (1944) axioms and time consistency. Traeger (2012) and Crost & Traeger (2010) argue in more detail why Epstein-Zin preferences are relevant to climate change evaluation.

We denote the measure of relative Arrow Pratt risk aversion by RRA, and the measure of aversion to intertemporal substitution, or the propensity for consumption smoothing, by $\eta$. In an intergenerational interpretation, the parameter $\eta$ can

[19] If we have a solution to the 4 state model for a related scenario, we obtain a very good initial guess for the 3 state problem by evaluating the 4 state value function at the final year of the planning horizon of the 3 state problem. Other guesses merely sum utility over a few centuries fixing the investment rate to a reasonable value and the abatement rate to 100%, independent of the initial state. A similar method iterates the stationary Bellman equation fixing exogenous parameters to their values at the end of the planning horizon (again replacing the processor intensive maximization step by merely guessing an optimal policy).

also be interpreted as the parameter of intergenerational consumption smoothing. Appendix A derives the Bellman equation disentangling risk aversion from intertemporal substitution

$$V^*(k_t, M_t, T_t, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta}\Delta t + \frac{\beta_{t,\Delta t}}{1-\eta}\left(\mathrm{E}\left[(1-\eta)V^*(k_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, t+\Delta t)\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}. \tag{22}$$

We solve the generalized Bellman equation (22) the same way as the original Bellman equation (19), and $V^*$ has the same normalization. Equation (22) uses a transformation explained in Traeger (2009) making the Bellman equation linear in the time step (as opposed to the original formulation of Epstein-Zin preferences). Under certainty, the non-linear exponents in equation (19) vanish and we are back in the setting of equation (19). The same observation holds if RRA $= \eta$, i.e., when risk preference happens to coincide with the decision maker's propensity to smooth consumption over time.

Vissing-Jørgensen & Attanasio (2003), Basal & Yaron (2004), Basal, Kiku & Yaron (2010), and Nakamura, Steinsson, Barro & Ursua (2010) provide preference estimates for Epstein-Zin preferences, explaining observed asset market behavior. These papers either estimate the preferences based on Campbell's (1996) approach of log-linearizing the Euler equation in the asset pricing context, or calibrate asset pricing models to the financial market data. A somewhat representative estimate is $\eta = \frac{2}{3}$ and a risk aversion coefficient around 10. Note that also the original DICE model picks $\eta$ based on observed market interest. With a single entangled parameter, or in a deterministic setting, however, the original DICE model cannot match both the risk-free interest (or discount) rate and the risk premia. The application of the model by Crost & Traeger (2010) analyzes the effects of preference disentanglement under damage uncertainty in detail, and Jensen & Traeger (n.d.) analyze the effects of preference disentanglement in the context of growth uncertainty.

# 4 Conclusions

We presented a stochastic dynamic programming model for analyzing uncertainty in the integrated assessment of climate change. In difference to Monte-Carlo studies, linearized models, or discretized models, we solve the full non-linear problem of finding optimal policies under uncertainty. The model fills the gap between highly stylized models treating uncertainty correctly, highly complex models that replace decision making under uncertainty by sensitivity studies, and a few state of the art high-dimensional stochastic integrated assessment models that require extremely powerful computer resources and/or pay the price of limiting the approximation quality and the uncertainty domain. The present model is a close relative to the wide-spread

DICE model, running in shorter time steps with an infinite planning horizon. We introduce Epstein-Zin preferences, which permit the modeler to distinguish between effects deriving from intertemporal consumption smoothing, and those stemming from risk and risk aversion.

We interpolated a complex set of exogenous difference equations in DICE by their continuous time solutions. We hope that our presentation of the slightly simplified DICE model makes this integrated assessment model even more accessible to the introspection of a large audience of environmental economists. Our model offers a low dimensional state space, without sacrificing much realism as compared to the full DICE model. The low dimensional state space has several important advantages when assessing optimal policies under uncertainty. First, the model solves on a regular PC in reasonable time. Second, it permits a significantly better approximation of the true value function solving the dynamic programming problem. Third, with given computer resources, we can increase the domain of the value function approximation and, thus, of the type and magnitude of uncertainty amenable to examination. Fourth, we can introduce (more) additional state variables that capture persistence in stochastic shocks, information and learning, and simultaneous uncertainties. We hope that the present model assists the community in pushing a more comprehensive, theoretically sound analysis of uncertainty in the integrated assessment of climate change.

# References

Basal, R., Kiku, D. & Yaron, A. (2010), 'Long run risks, the macroeconomy, and asset prices', *American Economic Review: Papers & Proceedings* **100**, 542–546.

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.

Campbell, J. Y. (1996), 'Understanding risk and return', *The Journal of Political Economy* **104**(2), 298–345.

Crost, B. & Traeger, C. (2010), 'Risk and aversion in the integrated assessment of climate change', *CUDARE Working Paper 1104* .

Epstein, L. G. & Zin, S. E. (1989), 'Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework', *Econometrica* **57**(4), 937–69.

Jensen, S. & Traeger, C. (n.d.), 'Growth uncertainty in the integrated assessment of climate change'.

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

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

Kelly, D. L. & Kolstad, C. D. (2001), 'Solving infinite horizon growth models with an environmental sector', *Computational Economics* **18**, 217–231.

Leach, A. J. (2007), 'The climate change learning curve', *Journal of Economic Dynamics and Control* **31**, 1728–1752.

Lemoine, D. & Traeger, C. (2011), 'Tipping points and ambiguity in the economics of climate change', *NBER Working Paper No. 18230* .

Miranda, M. & Fackler, P. L., eds (2002), *Applied Computational Economics and Finance*, Massachusetts Institute of Technology, Cambridge.

Nakamura, E., Steinsson, J., Barro, R. & Ursua, J. (2010), 'Crises and recoveries in an empirical model of consumption disasters', *NBER* **15920**.

Nordhaus, W. (2008), *A Question of Balance: Economic Modeling of Global Warming*, Yale University Press, New Haven. Online preprint: A Question of Balance: Weighing the Options on Global Warming Policies.

Traeger, C. (2009), 'Recent developments in the intertemporal modeling of uncertainty', *ARRE* **1**, 261–85.

Traeger, C. (2010), 'Intertemporal risk aversion – or – wouldn't it be nice to know whether Robinson is risk averse?'. CUDARE Working Paper No. 1102.

Traeger, C. (2012), 'Why uncertainty matters - discounting under intertemporal risk aversion and ambiguity', *CESifo Working Paper No. 3727* .

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

von Neumann, J. & Morgenstern, O. (1944), *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton.

Weil, P. (1990), 'Nonexpected utility in macroeconomics', *The Quarterly Journal of Economics* **105**(1), 29–42.

# Appendix

## A Normalizing the Bellman equation

The general Bellman equation expressed in terms of the original value function $V$ is

$$V(K_t, M_t, T_t, t) = \max_{C_t, \mu_t} L_t \frac{\left(C_t/L_t\right)^{1-\eta}}{1-\eta}\Delta t \; + \frac{\exp[-\delta_u \Delta t]}{1-\eta} \times \left(\mathrm{E}\;\left[(1-\eta)V(K_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, t+\Delta t)\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}.$$

Using effective labor units, i.e., the definitions $c_t = \frac{C_t}{A_t L_t}$ and $k_t = \frac{K_t}{A_t L_t}$, we bring this equation to the form

$$V(k_t A_t L_t, M_t, T_T, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} A_t^{1-\eta} L_t \Delta t + \frac{\exp[-\delta_u \Delta t]}{1-\eta} \exp[(1-\eta) g_{A,t} \Delta t] \times$$
$$A_t^{1-\eta} \exp[g_{L,t}\Delta t] L_t \left(\mathrm{E}\left[(1-\eta)\frac{V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{A_{t+1}^{1-\eta} L_{t+1}}\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}.$$

The second line inserted $A_{t+1}^{1-\eta} L_{t+1}$ in the denominator of the inner bracket, which cancels with the term $A_{t+1}^{1-\eta} L_{t+1} = \exp[(1-\eta) g_{A,t} \Delta t] A_t^{1-\eta} \exp[g_{L,t}\Delta t] L_t$ inserted outside of the brackets. We use the definition $V(K_t, M_t, t)|_{K_t = k_t A_t L_t} = V^*(k_\tau, M_\tau, \tau) A_t^{1-\eta} L_t$ and cancel $A_t^{1-\eta} L_t$ on both sides of the equation (and inside the square brackets for $t+1$), which delivers

$$V^*(k_t, M_t, T_t, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \frac{\exp[-\delta_u + g_{A,t}(1-\eta) + g_{L,t}]}{1-\eta} \times \left(\mathrm{E}\left[(1-\eta)V^*(k_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, t+\Delta t)\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}.$$

and, thus, equation (22) and its special case equation (19).

## B Calibration

Table B summarizes the model parameters. We calibrate our model such that the optimal time paths of four variables are similar to those predicted by the DICE-2007 model: $CO_2$ concentration, temperature, abatement rate, and social cost of carbon.[20]

[20] We used the EXCEL version downloadable from http://nordhaus.econ.yale.edu/DICE2007.htm to generate the optimal time paths of DICE-2007. It generates a longer time series than depicted for example in Nordhaus (2008). Note that the EXCEL model assumes a constant savings rate. We

Figures 2 and 3 show our calibration of equation (10) approximating the carbon cycle. Our carbon removal rate decreases from the initial value $\delta_{M,0}$ to the asymptotic value $\delta_{M,\infty}$, where the rate of decline is characterized by $\delta_M^*$. Figure 4 shows the calibration of the heat capacity and feedback related delay parameter $\sigma_{forc}$, and for the parameter $\sigma_{ocean}$ capturing ocean temperature related feedbacks (equation 18). Any individual parameter change improving the fit in one dimension worsens the fit in at least one of the other variables. Finally, Figure 5 depicts the difference between oceanic and atmospheric temperatures in DICE-2007, and compares it to our simple quadratic approximation stated in equation (12).

We solve the Bellman equations (19) or (22) by help of the function iteration algorithm described in section 3.2. For this purpose, we approximate the value function $V^*$ by Chebyshev polynomials. We update the coefficients by collocation at the Chebychef nodes spelled out in Table 2 (rectangular grid). To arrive at this final node grid, we sequentially increased the number of nodes in each dimension, until a further increase in the number of basis function no longer affected the solution. Figure 6 shows that a further increase of the node number beyond our $18 \times 6 \times 10 \times 6 = 6480$ nodes has no observable effect on increasing the accuracy of our simulation. Our convergence criterion was a coefficient change of less than $10^{-4}$. The corresponding maximal relative change in the value function was less than $10^{-10}$. Figure 7 shows that a further reduction of the convergence tolerance by an order of magnitude had no effect on the optimal time paths of the variables of interest.

did find an almost constant savings rate in our optimizing model as well, and the EXCEL version of DICE seems to be a close fit to the fully optimizing GAMS version as well for the time span for which we have both data series.

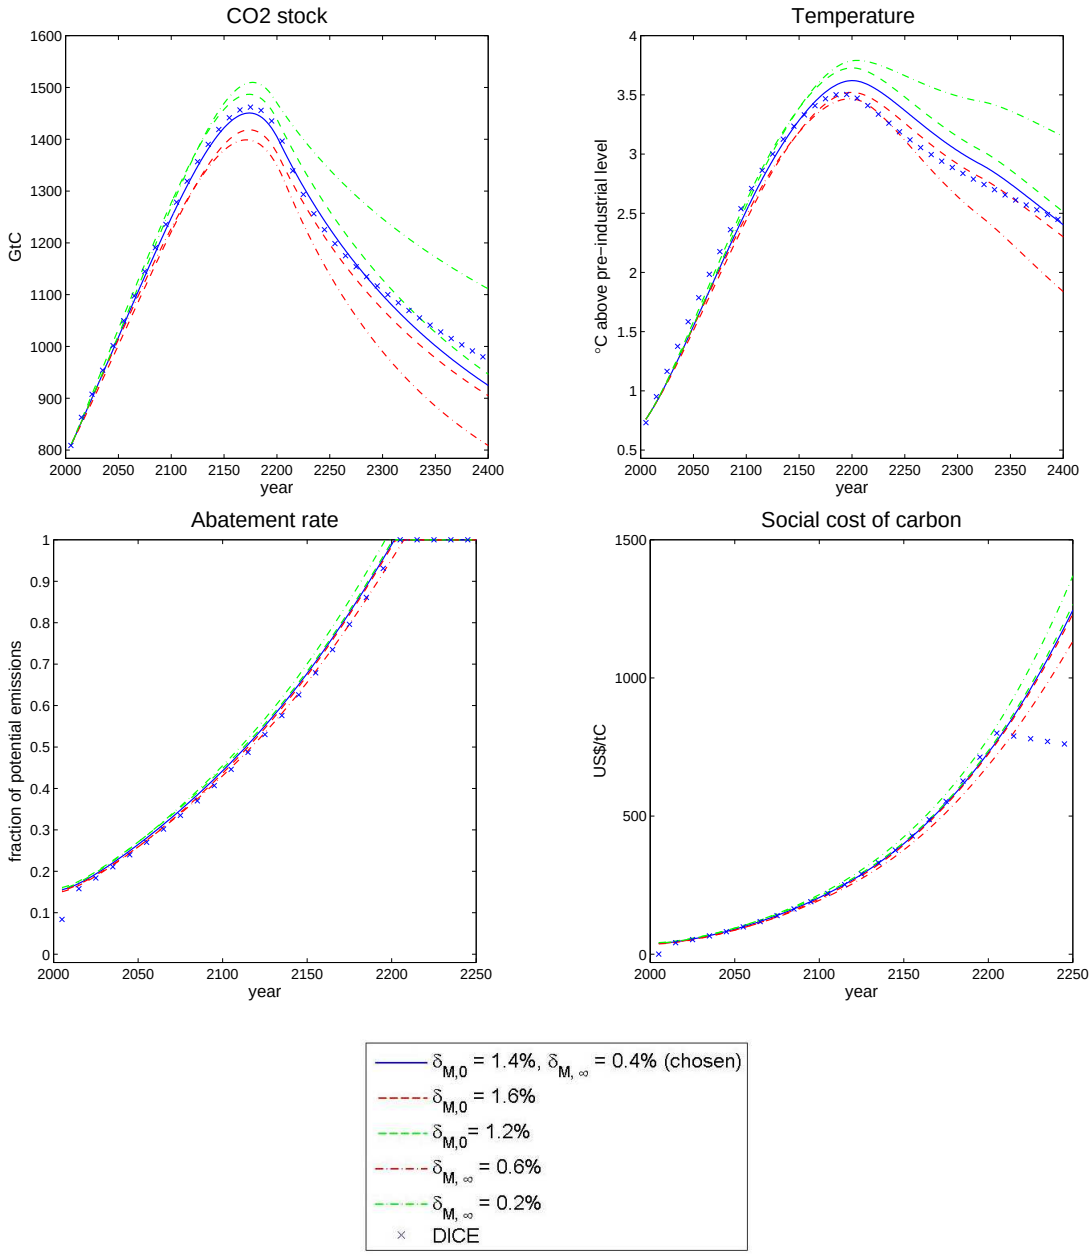


Figure 2 shows the calibration of the rate governing $CO_2$ removal from the atmosphere. We calibrate the initial rate $\delta_{M,0}$ and asymptotic rate $\delta_{M,\infty}$. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

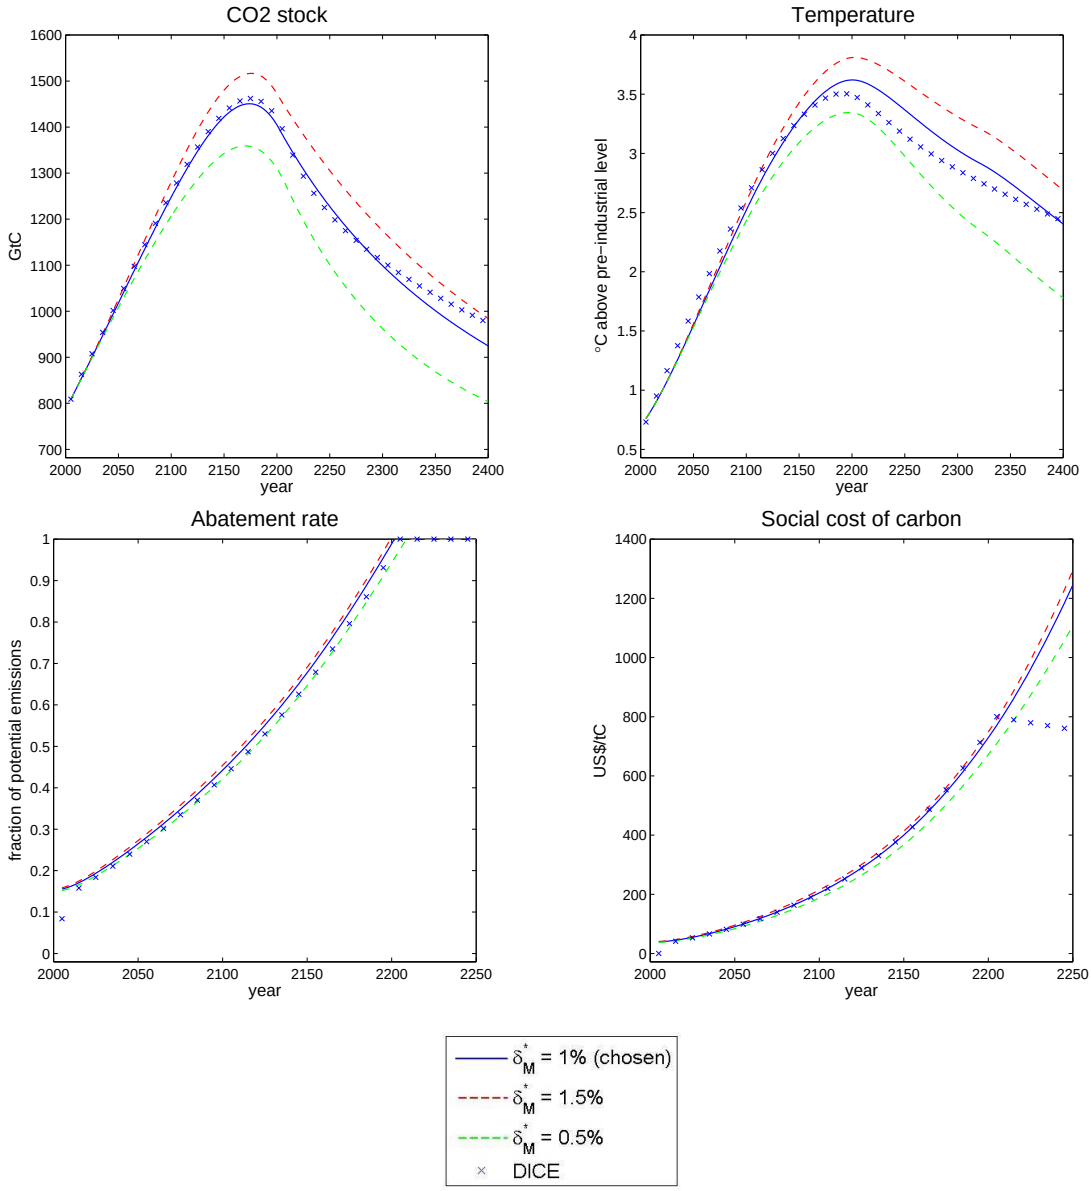


Figure 3 shows the calibration of the rate governing $CO_2$ removal from the atmosphere. We calibrate the parameter $\delta_M^*$ governing the speed of convergence from the initial to the asymptotic rate of $CO_2$ removal.

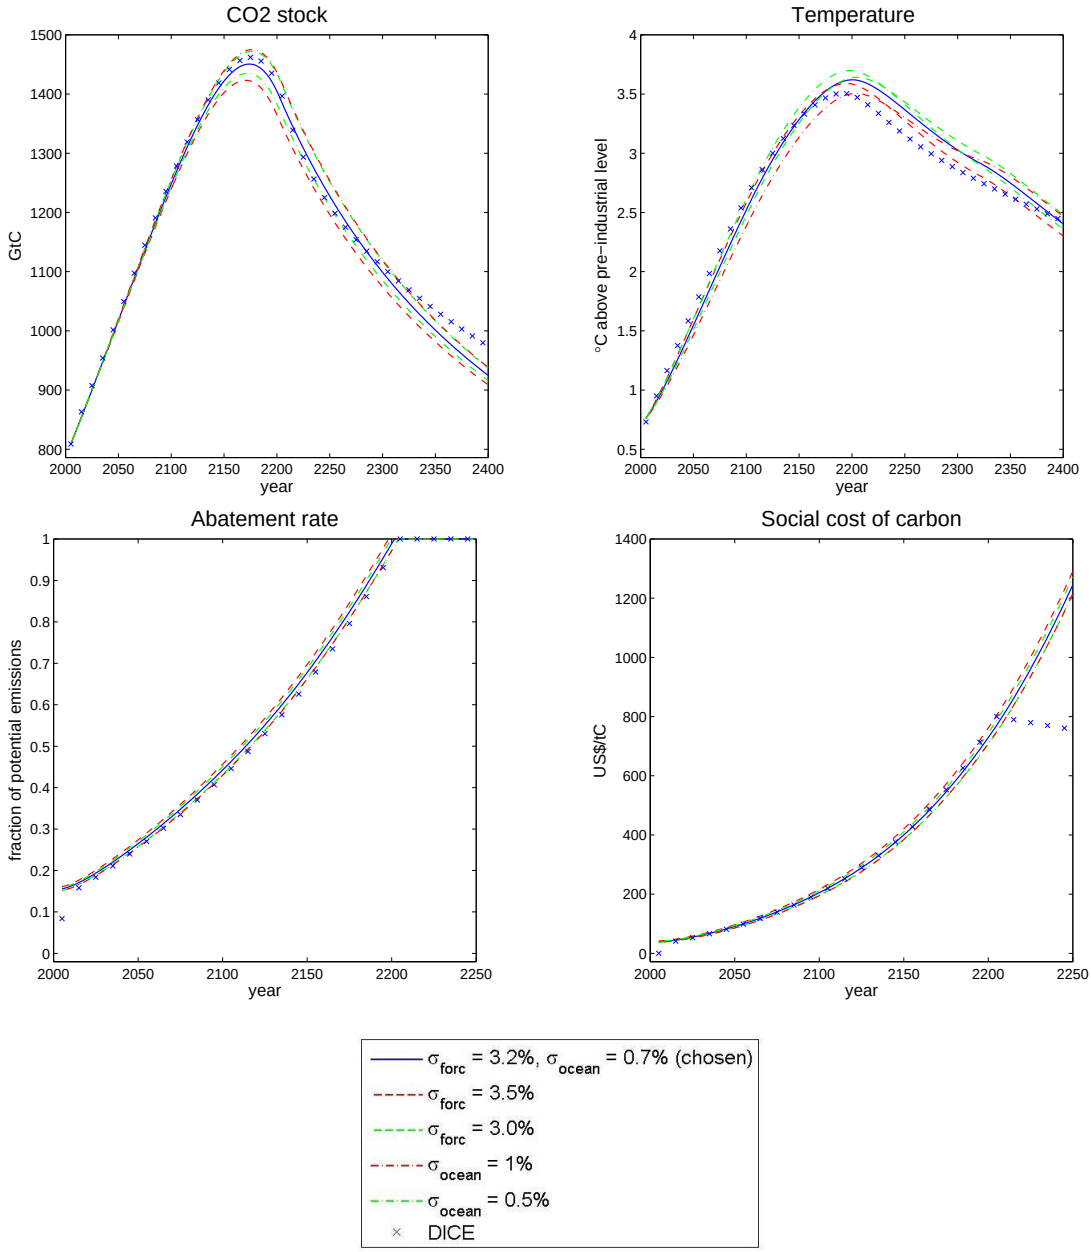


Figure 4 shows the calibration of the warming delay parameter $\sigma_{forc}$ and the parameter $\sigma_{ocean}$ connecting atmospheric and oceanic temperatures. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

Table 1 Parameters of the model

| | | |
|---|---|---|
| Economic Parameters | | |
| $\eta$ | 2 | intertemporal consumption smoothing preference |
| RRA | 2 | coefficient of relative Arrow-Pratt risk aversion |
| $b_1$ | 0.284% | damage coefficient |
| $b_2$ | 2 | damage exponent |
| $\delta_u$ | 1.5% | pure rate of time preference |
| $L_0$ | 6514 | in millions, population in 2005 |
| $L_\infty$ | 8600 | in millions, asymptotic population |
| $g_L^*$ | 3.5 | rate of convergence to asymptotic population |
| $K_0$ | 137 | in trillion 2005-USD, initial global capital stock |
| $\delta_K$ | 10% | depreciation rate of capital |
| $\kappa$ | 0.3 | capital elasticity in production |
| $A_0$ | 0.0058 | initial labor productivity; corresponds to total factor productivity of 0.02722 used in DICE |
| $g_{A,0}$ | 1.31% | initial growth rate of labor productivity; corresponds to total factor productivity of 0.9% used in DICE |
| $\delta_A$ | 0.1% | rate of decline of productivity growth rate |
| $\sigma_0$ | 0.1342 | $CO_2$ emissions per unit of output in 2005 |
| $g_{\sigma,0}$ | $-0.73\%$ | initial rate of decarbonization |
| $\delta_\sigma$ | 0.3% | rate of decline of the rate of decarbonization |
| $a_0$ | 1.17 | cost of backstop in 2005 |
| $a_1$ | 2 | ratio of initial over final backstop cost |
| $a_2$ | 2.8 | cost exponent |
| $g_\Psi^*$ | $-0.5\%$ | rate of convergence from initial to final backstop cost |
| Climatic Parameters | | |
| $T_0$ | 0.76 | in °C, temperature increase of preindustrial in 2005 |
| $M_{preind}$ | 596 | in GtC, preindustiral stock of CO2 in the atmosphere |
| $M_0$ | 808.9 | in GtC, stock of atmospheric $CO_2$ in 2005 |
| $\delta_{M,0}$ | 1.4% | initial rate of decay of CO2 in atmosphere |
| $\delta_{M,\infty}$ | 0.4% | asymptotic rate of decay of CO2 in atmosphere |
| $\delta_M^*$ | 1% | rate of convergence to asymptotic decay rate of CO2 |
| $B_0$ | 1.1 | in GtC, initial CO2 emissions from LUCF |
| $\delta_B$ | 1.05% | growth rate of CO2 emisison from LUCF |
| $s$ | 3.08 | climate sensitivity (equilibrium temperature response to doubling of atmospheric $CO_2$ concentration w.r.t. preindustrial) |
| $\eta_{forc}$ | 3.8 | forcing of $CO_2$-doubling |
| $\lambda$ | $\eta_{forc}/s \approx 1.23$ | ratio of forcing to temperature increase under $CO_2$-doubling |
| $EF_0$ | $-0.06$ | external forcing in year 2000 |
| $EF_{100}$ | 0.3 | external forcing in year 2100 and beyond |
| $\sigma_{forc}$ | 3.2% | warming delay, heat capacity atmosphere |
| $\tilde{\sigma}_{ocean}$ | 0.7% | parameter governing oceanic temperature feedback |

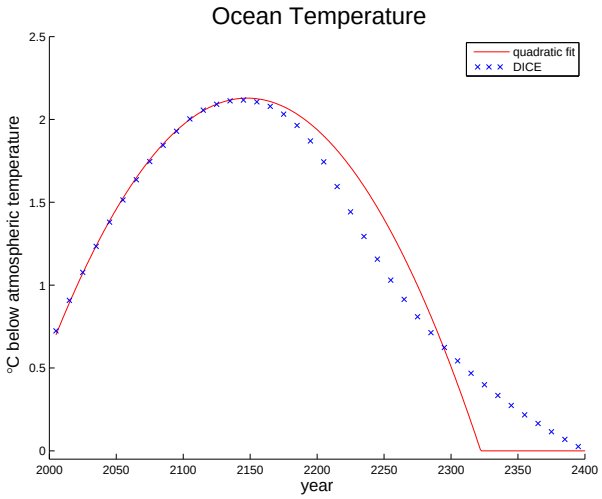


Figure 5 compares our simple quadratic approximation of the temperature difference between oceans and atmosphere in equation (12) to the actual difference resulting from the DICE-2007 model. The noticeable difference emerging two centuries into the future causes also the slightly more pronounced difference between our and DICE's atmospheric temperature observed after the year 200 in the earlier calibration plots. Section 2.3 notes a less simplistic approximation using a Chebychev interpolation, if desired.

Table 2 Location of Collocation Nodes

| Node | Effective Capital (k) | Carbon Stock (M) | Transformed Time ($\tau$) | Temperature (T) |
|---|---|---|---|---|
| 1 | 0.53 | 575 | 0.006 | 0.07 |
| 2 | 0.75 | 762 | 0.054 | 0.59 |
| 3 | 1.18 | 1087 | 0.146 | 1.48 |
| 4 | 1.81 | 1463 | 0.273 | 2.52 |
| 5 | 2.62 | 1788 | 0.422 | 3.41 |
| 6 | 3.59 | 1975 | 0.578 | 3.93 |
| 7 | 4.69 | | 0.727 | |
| 8 | 5.87 | | 0.854 | |
| 9 | 7.12 | | 0.946 | |
| 10 | 8.38 | | 0.994 | |
| 11 | 9.63 | | | |
| 12 | 10.81 | | | |
| 13 | 11.91 | | | |
| 14 | 12.88 | | | |
| 15 | 13.69 | | | |
| 16 | 14.32 | | | |
| 17 | 14.75 | | | |
| 18 | 14.97 | | | |

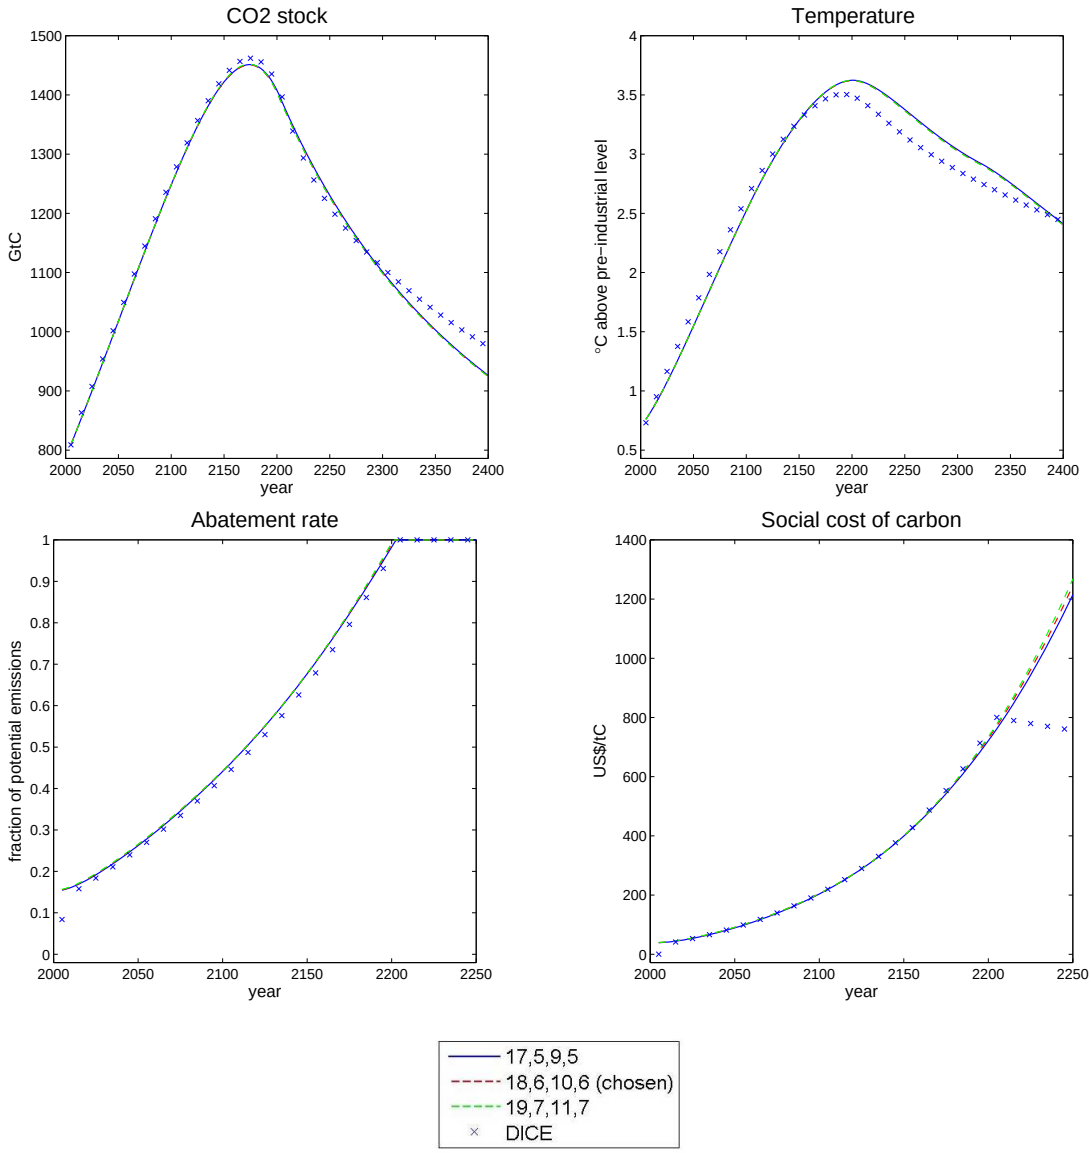


Figure 6 shows robustness of the results to variations of the number of basis functions (and corresponding collocation nodes). We use a tensor basis, i.e./ the basis containing all combinations of the first 18 Chebychev polynomials in the (effective) capital dimension, the first 6 Chebychev polynomials in the carbon stock dimension, the first 10 Chebychev polynomials in the time dimension, and the first 6 Chebychev polynomials in the temperature dimension.

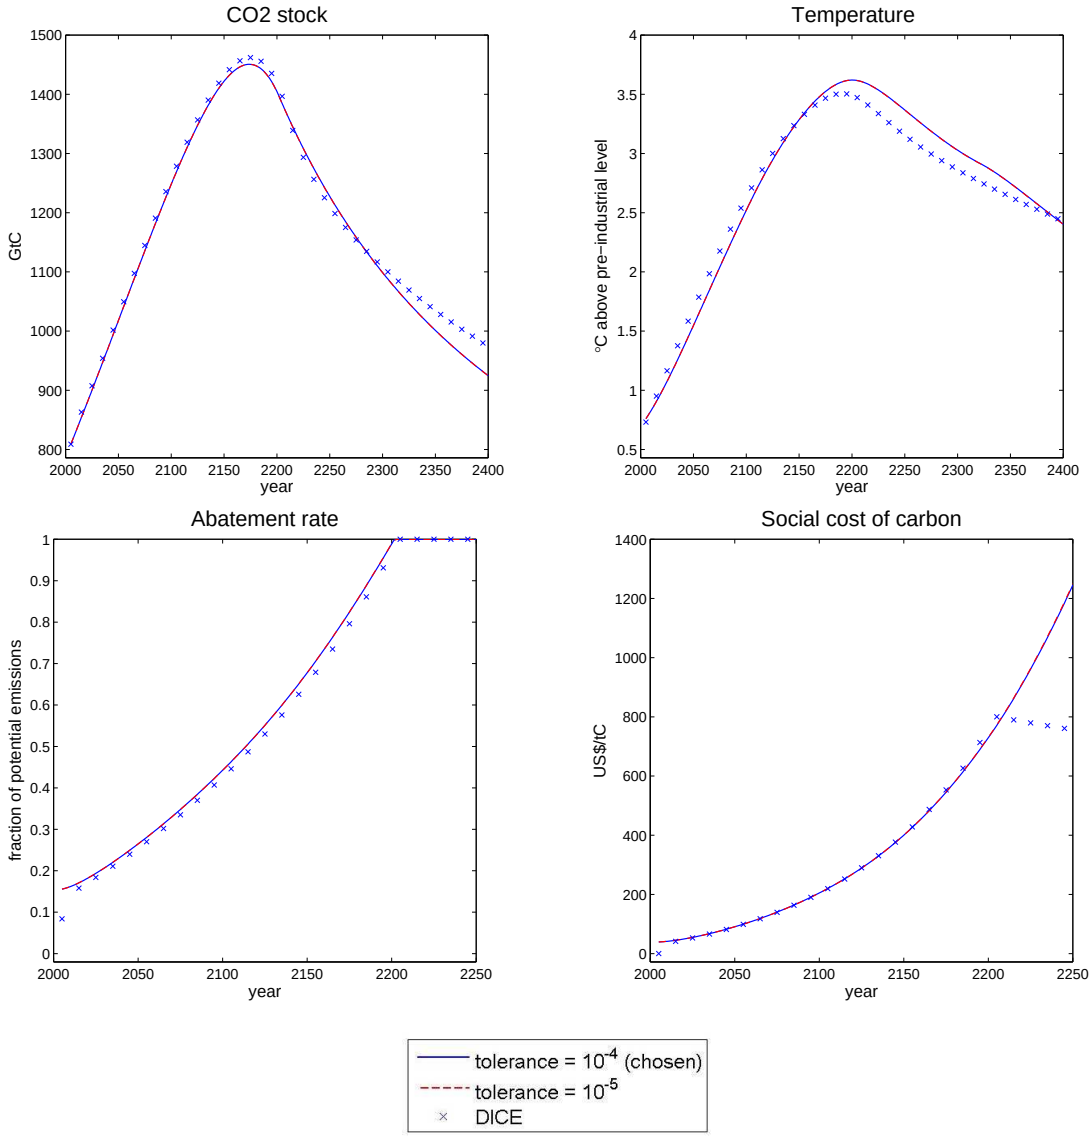


Figure 7 shows robustness of the results to a decrease in the convergence tolerance.