University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

Year 2013 *Paper 1130R*

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. Climate change, long-term economic development, and their interactions are highly uncertain. The quantitative analysis of optimal mitigation policy under 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 or policy function approximations, and reductions of the uncertainty domain. The paper promotes a state reduced, recursive dynamic programming implementation of the DICE-2007 model. We achieve the reduction by simplifying the carbon cycle and the temperature delay equations. We compare our model's performance and that of the DICE model to the scientific AOGCM models emulated by MAGICC 6.0 and find that our simplified model performs equally well as the original DICE model. Our implementation solves the infinite planning horizon problem in an arbitrary time step. The paper is the first to carefully analyze the quality of the value function approximation using two different types of basis functions and systematically varying the dimension of the basis. We present the closed form, continuous time approximation to the exogenous (discretely and inductively defined) processes in DICE, and we present a numerically more efficient re-normalized Bellman equation that, in addition, can disentangle 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, MAGICC, basis, 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. Most integrated assessment models (IAMs) are too complex to permit a proper incorporation of uncertainty. This limitation weighs strongly because optimal policy today depends on forecasting the climate, the global economy, and their interactions over at least a couple of centuries. The present paper reduces the complexity of a wide-spread IAM, DICE, by simplifying the climate's representation, without sacrificing the original model's quality in reproducing the relation between emissions and temperature change found in the "big" climate science models, the Atmosphere-Ocean Global Circulation Models (AOGCMs).

Nordhaus's (2008) DICE model combines a Ramsey-Cass-Koopmans growth economy with a simple model of the carbon cycle and the climate. DICE is one of the most wide-spread integrated assessment models. Three advantages of DICE are that, 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. It therefore has a large audience familiar with the basics of the model. Third, the US government employs the DICE model, in combination with FUND and PAGE, to determine the federal social cost of carbon (Interagency Working Group on Social Cost of Carbon 2010, Interagency Working Group on Social Cost of Carbon 2013).

DICE's parsimony implies an original state space reducible to six state variables (plus time). To implement stochastic shocks, persistent uncertainty, or Bayesian learning, we have to implement DICE as a recursive dynamic programming model. The curse of dimensionality in dynamic programming makes six (or seven) state variables a very large model. The cost of increasing the dimension is a combination of processor time, sacrifices in the numerical approximation to the true solution, and/or introducing uncertainty only as small fluctuations around the deterministic path. The most interesting uncertainties for climate change policy imply future states of the world that can vary largely with respect to their expected values. Moreover, persistence in shocks (e.g. on economic growth) or Bayesian learning (e.g. over climate sensitivity) require additional informational state variables. Similarly, extensions on the economic side of DICE (e.g. improving the abatement cost dynamics) require additional state variables. The current model reduces the climate's representation by 3 states and, thus, cuts DICE's state space into half (plus time). It thereby permits a combination of more informational states (uncertainty), extending the economic dynamics, and improving numerical accuracy while allowing the uncertain variables to stray sufficiently far from their expected values.

I briefly summarize how our model differs from the original DICE-2007 model. 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 arbi-

trary time step without recalibrating the exogenous processes. Second, we implement a flexible time step. In particular, we suggest running the model in a annual rather than decadal time step when analyzing economic fluctuations and shocks, and when modeling learning over the climate system. Third, we simplify the ocean heat uptake related temperature delay, approximating the evolution of atmospheric temperature in a single delay equation (cutting one state). Fourth, we approximates the carbon cycle by means of a time (or more generally state) dependent rate of atmospheric carbon removal. Fifth, we normalize the model to effective labor units. This step reduces the node density required to achieve a given precision in the approximation and ensures that the infinite planning horizon problem converges on a feasible, compact support. Sixth, we present the dynamic programming equation for Epstein-Zin-Weil preferences. 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 variable. Epstein-Zin-Weil 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 an application of the current model, Crost & Traeger (n.d.) demonstrate the relevance of Epstein-Zin preferences when evaluating climate policies with DICE.

Kelly & Kolstad (1999, 2001) implement the DICE-1994 model as a recursive dynamic programming model, analyzing uncertainty and learning over the sensitivity of temperatures with respect to carbon emissions. Leach (2007) implements the same DICE version showing that learning slows down further under additional uncertainty.[1] These papers are seminal contributions to uncertainty assessment in climate change and careful implementations of the original DICE version. A major update of this original version starting with DICE-1999 replaces a constant carbon decay rate by an explicit 3 box carbon cycle (Nordhaus & Boyer 2000). Other updates include the introduction of CO_2 emissions from land use change and forestry and time varying abatement cost.[2] Our recursive implementation of DICE is similar in spirit to Kelly & Kolstad (1999, 2001). With respect to their work, we update the model to DICE-2007[3], add a flexible time step, and re-normalize the model to effective labor units enabling convergence on a (reasonably) bounded capital interval for an infinite time horizon. The important extensions are threefold. First, instead of assuming a constant decay of atmospheric excess carbon as in DICE 19944, we imitate the recent DICE carbon cycle by means of a time (and possibly state) dependent rate of carbon

[1] Leach (2007) also cites Nordhaus & Boyer (2000), but the DICE-1999 model described there already contains a 3 box carbon cycle model.

[2] External forcing from non-CO_2 greenhouse gases is already part of the DICE-1994 model even though it was not taken up in the cited recursive implementations.

[3] The recently released DICE 2013 version mostly contains parametric but not structural updates of the model. We test our main contribution, the simplified climate part of the model as well against DICE-2013.

removal. Second, we simplify the warming delay equations to eliminate the explicit ocean temperature state. Third, we compare our model's and DICE's temperature response to those of the scientific AOGCMs for the emission scenarios used by the International Panel on Climate Chanve (IPCC). Fourth, we use a Chebychev basis for our value function approximation, compare it to the performance of a spline basis, and show what number of basis function is needed to reach full numerical precision, i.e., invariance of the results to further increases in the number of basis functions or reductions in tolerance.

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, paralellized 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 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 2013). Finally, the model relates to a set of more stylized stock pollution models analyzing in dynamic programming settings the optimal choice of the policy *instrument* (taxes versus quantities) to control emissions. These models employ linearizations for uncertainty assessment, linear-quadratic formulations from the start, or significantly more stylized models of climate and abatement (Hoel & Karp 2001, Hoel & Karp 2002, Kelly 2005, Karp & Zhang 2006, Heutel 2011, Fischer & Springborn 2011); these models are not designed to quantify the optimal abatement *level*.

Section 2 presents the continuous time approximation to the exogenous processes in DICE and the endogenous equations of motion for a flexible time step. Section 3 discusses our approximations to the carbon cycle and the warming delay equations. Section 4 derives the normalized Bellman equation, extends it to Epstein-Zin preferences, and summarizes the solution algorithm. Section 5 presents the result of our model calibration, and compares its temperature response to DICE and an average of scientific AOGCMs for the IPCC's emission scenarios. It also presents a comparison of results using a spline versus a Chebychev basis. Section 6 concludes. The appendix presents the details of deriving the normalized Bellman equation and the calibration, and discusses further comparisons between our model, DICE, and the AOGCM average.

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 discuss how uncertainty modifies the transition equations. 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 increases over time. In section 3, we discuss two additional, exogenous processes resulting from our carbon cycle and the warming delay approximation. We abbreviate growth rates by g and rates of decay by δ.[4]

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 in DICE have the continuous time approximation[5]

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

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_∞ 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

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

[4] A subindex indicates the growing or decaying variable. We use δ_A instead of δ_{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.

[5] 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 β_t in the renormalized Bellman equation as we explain in section 4.

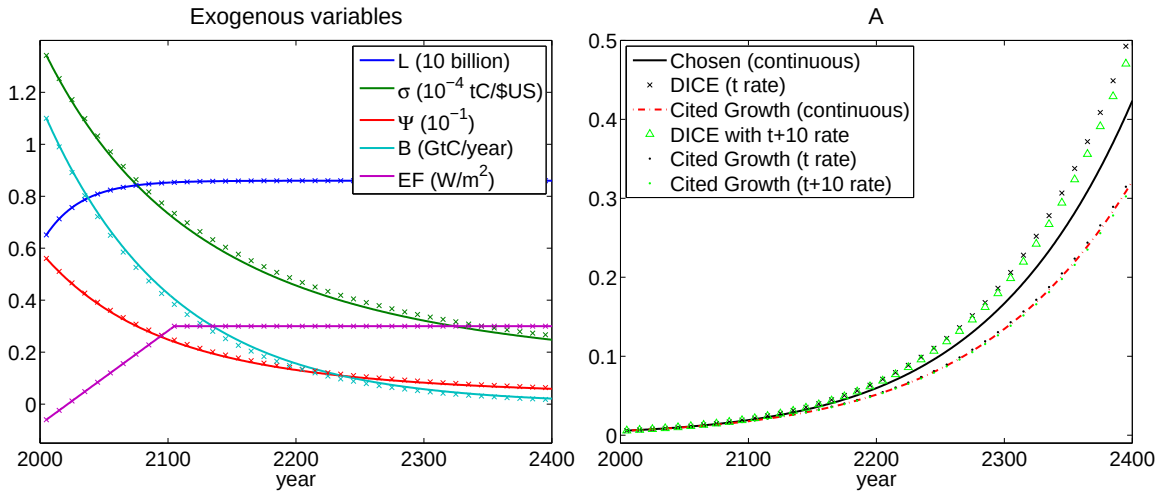

Figure 1 shows the time paths of the DICE model's exogenous parameters. The graph on the left shows our continuous time interpolations (solid) and the original, recursively defined values in the DICE model (crossed) for population (L), carbon intensity of production (σ), abatement cost coefficient (Ψ), 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 beginning of the according decade. Together, they bound a corresponding continuous model.

leading to the analytic continuous time solution

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

The parameter $g_{A,0}$ denotes the initial growth rate and δ_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. This reformulation improves the numerical approximation for a given number of basis functions in the capital dimension

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 the cited rate $g_{A,t}$. The dash-dotted line in Figure 1 depicts the continuous time productivity path that 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 employ $g^{DICE}_{A,0}/10$ 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 $g_{A,0}^{DICE}$ (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 coefficient Ψ_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 9 below). The rate g_Ψ^* 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

[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 $1-\kappa$ (κ denotes the factor share of capital, see equation (9) and the preceding paragraph). 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.

[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.

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.

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 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. The present section discusses the deterministic equations of motions replicating DICE-2007. The subsequent section discusses how to introduce different types of uncertainty into the model.

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 κ 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 $y_t^{gross} = \frac{Y_t^{gross}}{A_t L_t} = k_t^{\kappa}$. In effective labor units, the capital state is largely constant and, in the absence of technological uncertainty, allows us to represent capital on a tight state space using few nodes.[9] We introduce a flexible time step Δt, keeping our flow variables, including consumption and production, defined in units per year. Then, production during the

[9] In the usual growth model, non-normalized capital grows without bounds, leaving quickly any finite numerical support. The DICE-2007 model assumes falling growth rates of population and technology level. Thus, in principle, at some point in the very long run future capital converges, but at a level that would marginalize the evolution of capital over the next couple of centuries that are most relevant for present climate policy. The per effective labor normalization implies a well-defined infinite horizon limit on a narrow state space, allowing for a good resolution of both current and future deviations from some steady state level.

period $[t, t+\Delta t]$ is $k_t^\kappa \Delta t$. Our base case model calibration and results presented in section 5 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

$$y_t = \frac{1-\Lambda(\mu_t)}{1+D(T_t)}\ k_t^\kappa = \frac{1-\Psi_t \mu_t^{a_2}}{\left(1+b_1 T_t^{b_2}\right)}\ k_t^\kappa\ . \tag{9}$$

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

$$k_{t+\Delta t} = \left[(1-\delta_k)^{\Delta t}\, k_t + y_t \Delta t - c_t \Delta t\right] \exp[-(g_{A,t}+g_{L,t})\Delta t]\ , \tag{10}$$

where δ_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

$$E_t = (1-\mu_t)\,\sigma_t A_t L_t k_t^\kappa + B_t\ . \tag{11}$$

Industrial emissions are proportional to gross production $A_t L_t k_t^\kappa$, and the emission intensity of production σ, and they are reduced by the emission control rate μ_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

$$M_{t+\Delta t} = M_{pre} + (M_t - M_{pre})\,(1-\delta_{M,t})^{\Delta t} + E_t \Delta t\ . \tag{12}$$

The pre-industrial emission stock M_{pre} is the steady state level in the absence of anthropogenic emissions. Equation (12) is our approximation to the carbon cycle in

DICE-2007. Section 3 introduces two different calibrations for the (time or state-dependent) rate of carbon removal from the atmosphere $\delta_{M,t}$.

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 η_{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

$$T_{t+\Delta t} = (1 - \sigma_{forc})T_t + \sigma_{forc} \ \frac{F_{t+\Delta t}}{\lambda} - \sigma_{ocean} \ \Delta T_t \ . \tag{13}$$

The parameter

$$\sigma_{forc} = 1 - (1 - \sigma_{forc}^{ann})^{\Delta t} \tag{14}$$

captures the warming delay and, thus, scales with the model's time step. The parameter

$$\sigma_{ocean} = 1 - (1 - \sigma_{ocean}^{ann})^{\Delta t} \tag{15}$$

quantifies the ocean cooling in a given time step that derives from the atmosphere ocean temperature difference ΔT_t. This last term in equation (13) replaces the oceanic temperature state in DICE-2007. Section 3 discusses two alternative calibrations of the term ΔT_t, and the scaling of the parameters.

2.3 Uncertainty

The paper develops a model for analyzing uncertainty in the integrated assessment of climate change. This section explains how stochasticity, persistent shocks, and Bayesian learning alter the equations of the base model. We discuss these modifications in general, and refer to different applications.

Let $\mathbf{S}$ denote the vector of all state variables, and let S^o denote a generic state variable affected by uncertainty, e.g. the capital stock K, carbon concentration M, or temperature T. If we introduce persistent uncertainty over an originally exogenous variable like population L, or technology level A, or over model parameters (e.g. governing emission intensity or abatement costs), then we have to include these formerly exogenous variables as states as well. Let $\mathbf{P}$ denote the vector of model parameters governing the equations of motion.

The simplest form of introducing uncertainty are iid shocks on the right hand side of the equations of motion. Then, in addition to $\mathbf{S}$ and $\mathbf{P}$, the transition equation

$\tilde{S}^o_{t+\Delta t} = g(\mathbf{S}_t, \mathbf{P}, \tilde{\epsilon}_t)$ depends on the stochastic iid shock $\tilde{\epsilon}_t$. The tilde emphasizes uncertainty. For example, we could add additive noise in the carbon accumulation by transforming equation (12) to the form $\tilde{M}_{t+\Delta t} = M_{pre} + (M_t - M_{pre})(1 - \delta_{M,t}\Delta t) + E_t\Delta t + \tilde{\epsilon}_t$. If we model a persistent shock $\tilde{z}$, the shock itself becomes a new state variable governed by an equation of motion $\tilde{z}_{t+\Delta t} = h(z_t, \tilde{\epsilon}_t)$; and z_t becomes an additional state in the equation of motion for $S^o_{t+\Delta t} = g(\mathbf{S}_t, \mathbf{P}, z_t)$. Applications of models with growth shocks on the technology level in the climate policy context are Kelly (2005), Heutel (2011), Fischer & Springborn (2011), Cai, Judd & Lontzek (2012*b*), and an application with persistent shocks to technological progress generating long-term uncertainty is Jensen & Traeger (2013*a*). Note that the dynamic programming approach ensures that the decision maker always responds optimally to the realizations of the shocks. More importantly, the decision-maker in the present anticipates that the future decision makers will respond optimally to the actual evolution of the states.

The uncertainties characterized in the previous paragraph do not evolve over time: whereas the stochastic element might dependent on a persistent component, there is no structural learning. In some occasions, we would like to capture that future decision makers not only have a better knowledge of future states, but also gathered more structural information about the equations governing e.g. the climate system. Kelly & Kolstad (1999), Leach (2007), Kelly & Tan (2013), and Jensen & Traeger (2013*b*) discuss the policy implications of structural learning about climate sensitivity, i.e., the temperature response to the carbon stock. They treat climate sensitivity as an unknown parameter, and model Bayesian learning over its distribution. Bayesian learning combines a stochastic shock (which prevents immediate learning) with a subjective prior capturing uncertainty that the decision maker learns over time through observation. A tractable Bayesian learning model usually assumes that the stochasticity (likelihood) and the subjective uncertainty (prior) are conjugates, ensuring that the updated subjective prior belongs to the same class of probability distributions as the initial prior. Then, we characterize the prior by the parameters of the distribution $\mathbf{\Phi}$. These parameters become additional, informational state variables, and their equations of motion describe the evolution of learning $\mathbf{\Phi}_{t+\Delta t} = h(\mathbf{S}_t, \mathbf{P}_t, \mathbf{\Phi}_t)$. In the wide-spread example of a normal-normal learning model, prior and likelihood function interact additively, are both normally distributed, and the characteristic parameters of the distribution are the mean and the variance of the normal prior. To condense notation, the vector $\mathbf{\Phi}_t$ will henceforth capture exogenous variables that became endogenous because of uncertainty, informational states for Bayesian learning, and persistent components of stochastic shocks.

3 Discussion of the state reducing approximations

We replace the carbon cycle by a state dependent rate of carbon removal from the atmosphere. Similarly, we replace the endogenous evolution of ocean temperature by

a state dependent ocean cooling term. First, we discuss our qualitative simplifications with respect to the original DICE model. Second, we discuss our base case calibration of the state reduced climate equations. Here, carbon removal and ocean cooling only depend on time. Section 5.2 shows that this simple calibration competes well with the original DICE model in replicating the climate dynamics of the "big" scientific climate models (AOGCMs). Third, we discuss a more sophisticated interpolation of DICE's carbon cycle and temperature equations. Here, carbon removal and ocean cooling are also functions of the atmospheric carbon content. This more sophisticated approach replicates DICE's carbon and ocean cooling more closely for perturbations around the deterministically optimal path.

3.1 Qualitative discussion of the simplifications

In DICE 2007, the **carbon cycle** is a three box model. The first box, or carbon stock, characterizes atmospheric CO_2 (measured in terms of carbon content). This first atmospheric box absorbs the anthropogenic emissions. As atmospheric concentrations increase, atmospheric carbon diffuses into a second box, a carbon stock that jointly measures the carbon content of the shallow ocean and the biosphere. Finally, carbon from this second box diffuses into a third box, the deep ocean. Once anthropogenic emissions decline, the equilibrium net flow resulting from the diffusion can turn around, releasing carbon from the shallow ocean into the atmosphere. Nordhaus calibrates the diffusion parameters to simulations of the MAGICC model that we briefly discuss in section 5.2. The simplified three box model captures two important effects of carbon removal from the atmosphere. First, if antropogenic emissions increase, so does the partial pressure of atmospheric carbon relative to the shallow ocean. In consequence, we observe a higher net flow of carbon into the ocean. Second, in the long-run, the shallow and the deep oceans fill up with carbon and reduce their uptake. Then, less carbon leaves the atmosphere into these natural carbon sinks.

We replace DICE's three box model of the carbon cycle by an exogenous **rate of removal of excess carbon from the atmosphere** $\delta_{\mathbf{M,t}}$. For simplicity, we frequently call it a decay rate, acknowledging that atmospheric CO_2 does not actually decay, but only moves into the biosphere and shallow ocean. Our decay rate only acts on the carbon in excess of the long-term equilibrium concentration preceding the industrial revolution. Section 5.2 shows that our simplified climate system competes well with the original DICE model. Here we discuss why replacing DICE's carbon cycle with a time-dependent but exogenous rate of carbon removal only has a second or third order policy impact. Uncertainty causes a first order deviation of the optimal 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. Simulations with the original DICE model show that even largely different abatement scenarios imply almost the same decay rate of excess carbon for 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. This second order change in the decay rate is a flow change and, thus, only builds over time into a notable stock (concentration) difference. The resulting difference in the atmospheric CO_2 concentration then impacts warming. But warming is governed by a delay equation, delaying 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. The approximation *eliminates two state variables* from the original DICE model, corresponding to the carbon stock in the deep ocean, and the combined carbon content of the shallow ocean and biosphere.

Our model also simplifies the **temperature delay dynamics** with respect to the original DICE formulation. A given increase in the atmospheric CO_2 concentration results in a corresponding change of the radiative forcing, which is a measure for the change in the plant's energy balance. We can think of radiative forcing as the flame that we turn on (or up) to warm a big pot of soup. The radiative forcing does not immediately warm the planet to a new equilibrium temperature, but only step by step from period to period. The temperature in the next period is a function of the temperature the current period and the radiative forcing. The DICE model captures the major part of the warming delay by making atmospheric temperature a delay equation

$$T_{t+10} = T_t + C1 \left[F_{t+10} - \lambda T_t + C3(T_t^{Ocean} - T_t) \right].$$

The parameter λ specifies the equilibrium relation between anthropogenic forcing F and temperature change T, and the parameters $C1, C3$, and $C4$ (below) are constants in the original DICE model. In addition, the oceans with temperature T_t^{Ocean} keep cooling the atmosphere until their temperature catches up with the atmospheric temperature increase. The ocean temperature itself is modeled as another delay equation

$$T_{t+10}^{Ocean} = T_t^{Ocean} + C4(T_t - T_t^{Ocean}) = (1 - C4)T_t^{Ocean} + C4\, T_t.$$

Note that energy balance based climate models like MAGICC, to which DICE is calibrated, generate the warming delay almost entirely through explicit models of oceanic heat uptake. However, for a state reduced model, the delay equations employed by Nordhaus are significantly more convenient. In fact, we find that we can completely eliminate the ocean temperature as a state variable, and still achieving a similarly good approximation to the actual temperature dynamics.

We model oceanic cooling of the atmosphere through the **atmosphere-ocean temperature difference** $\mathbf{\Delta T_t} = T_t - T_t^{Ocean}$. We rewrite the first of the two equations above as

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

and replace the second by a time dependent, and possibly carbon stock dependent, approximation of the atmosphere-ocean temperature difference ΔT_t. 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. The subsequent sections explain the downscaling of the temperature delay equation to smaller time steps and two alternative calibrations of the atmosphere-ocean temperature difference ΔT_t.

3.2 The simple base case calibration

We now discuss how we calibrate equation (12) replacing the carbon cycle, and equation (13) replacing the ocean temperature delay dynamics. The rate $\delta_{M,t}$ of atmospheric CO_2 removal (in excess of preindustrial levels) governs the carbon dynamics in equation (12). In our first calibration, we make this decay rate a parametric function of time

$$\delta_{M,t}(t) = \delta_{M,\infty} + (\delta_{M,0} - \delta_{M,\infty}) \exp[-\delta_M^* \; t] \; . \tag{16}$$

The parameter $\delta_{M,0}$ characterizes the initial decay of excess carbon, the parameter $\delta_{M,\infty}$ captures the long-run decay rate, and the parameter δ_M^* characterizes the speed of change from the initial decay rate to the long-run decay rate. Figures 5 and 6 in Appendix B show our calibration.

The temperature equation (13) depends on the the parameter σ_{forc} governing warming delay, the parameter σ_{ocean} governing ocean cooling, and the atmosphere-ocean temperature differential ΔT_t. The calibration in the next section presents a sophisticated approach to reproducing the temperature differential along the optimal path and its perturbations. The present calibration seeks a simple parametric form that represents the temperature difference between the atmosphere and the ocean reasonably well along the optimal policy trajectory of the original DICE model. Having tried a set of simple parameteric functional forms, we decided for the following combination of the max and a linear-quadratic approximation

$$\Delta T_t(t) = \max\{0.7 + 0.02t - 0.00007t^2, 0\} \; . \tag{17}$$

Figure 8 in Appendix B shows that equation (17) captures the atmosphere-ocean temperature difference ΔT_t of the original model almost perfectly for the first 150 years. Later, it deviates notably until both the true temperature difference and its approximation converge to zero in the equilibrium. Our model calibrations in the appendix (e.g. Figure 5) reflect a small deviation of our atmospheric temperature path from that 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 because of discounting.

Section 5.1 laid out the relation between our equation (13) for the temperature dynamics and the corresponding system of equations in the original DICE model for its decadal time step. An exact mapping between a decadal and our finer time resolution is only possible for specific or stationary trajectories. The next section pursues a mechanical downscaling replicating the temperature dynamics in a stationary world. Here, we view the parameters σ_{ocean} and σ_{forc} as free parameters and calibrate them for a one year time step model to replicate the original DICE model's decadal output. Figure 7 in Appendix B shows our calibration. It turns our that the delay parameter σ_{forc} calibrates to the same parameter predicted by the mechanical downscaling in the subsequent section, whereas we find a marginally lower ocean cooling parameter σ_{ocean} in the present approach.

3.3 The sophisticated calibration

A more sophisticated DICE interpolation proceeds as follows. We first determine the deterministically optimal trajectory, and then run the original DICE model for variations of the optimal emissions path. These emission responses of the original DICE model result in an average time trend and a residual that varies across trajectories. We fit the time trend using a set of basis functions and regress the residual on one of the (other) state variables in the model, e.g., on the atmospheric carbon stock or the temperature.

We now proceed to explain our more sophisticated interpolation of the full DICE model. Using the original DICE model, we determine the deterministically optimal emission trajectory E_t^*, $t > 0$. We then run the original model in non-optimization mode, feeding variations of the emission trajectory. In the example discussed here, whose output we present in section 5.3, we vary emissions from half the optimal emission level to twice the optimal level in steps of one quarter, generating seven trajectories $E_t^i = iE_t^*$ for all t and $i \in \{0.5, 0.75, 1, 1.25, 1.5, 1.75, 2\}$. We then calculate the rate of carbon removal from the atmosphere $\delta_{M,t}$, and the temperature difference ΔT_t between the atmosphere and ocean along these seven trajectories. These trajectories result in an average time trend that we fit as $\delta_{M,t}^\dagger(t)$ and $\Delta T_t^\dagger(t)$ using flexible functional forms, here a third order spline with 30 nodes.

We subsequently regress the seven detrended residuals in every time period (decade) on the carbon stock. We thereby obtain linear, period-specific corrections of the general time-trend of the form $\Delta T_a(M) = n_a + m_a M_a$, where a is a particular decade and $n_a, m_a \in \mathbb{R}$ are the period-specific regressions parameters. For the function iteration, we have to evaluate the decay rate and the atmosphere-ocean temperature difference at arbitrary points in the continuous time state space.[10] For this purpose, we take a time-distance weighted, convex combination of the decay rate and temperature dif-

[10]For generating time paths we need to evaluate decay and temperature difference on the smaller time step, here, annual. But, already in the function iteration, the numerically optimal grid, e.g. consisting of Chebychev nodes, will not coincide with the decadal information deriving from the original DICE 2007 model.

ference at the nearest neighbors. Let the original model generate values $\Delta T_a(M)$ and $\Delta T_b(M)$, on a decadal scale at points a and $b = a + 10$. We then construct the value at time t, $a < t < b$, by setting $\Delta T_t^M(t, M) = \frac{b-t}{10}\Delta T_a(M) + \frac{t-a}{10}\Delta T_b(M)$. Following the same procedure, we generate the carbon stock based correction of the decay rate as $\delta_{M,t}^M(t, M) = \frac{b-t}{10}\delta_{M,a}(M) + \frac{t-a}{10}\delta_{M,b}(M)$.

The total rate of removal of excess carbon in the atmosphere is then the sum of the rate captured by the time trend $\delta_{M,t}^\dagger(t)$ and the period-specific, linear, carbon stock based correction $\delta_{M,t}^M(t, M)$

$$\delta_{M,t}^{dec}(t, M) = \delta_{M,t}^\dagger(t) + \delta_{M,t}^M(t, M) \ . \tag{18}$$

Similarly, we obtain the overall temperature difference as the sum of the time trend $\Delta T_t^\dagger(t)$ and the period-specific, linear, carbon stock based correction $\Delta T_t^M(t, M)$

$$\Delta T_t(t, M) = \Delta T_t^\dagger(t) + \Delta T_t^M(t, M) \ . \tag{19}$$

Figure 9 shows that the flexible time trend together with the simple, linear, carbon stock based correction nicely reproduces both the rate of atmospheric carbon removal and the temperature differences for the original perturbation scenarios.

Finally, we have to rescale the time step of the equations. A one to one mapping of the general non-stationary dynamics between a decadal and a finer time step is not possible. Here, we assume constancy of forcing, feedbacks, and emissions in establishing the equivalent dynamics for different time steps. E.g., the decadal equation for atmospheric temperature in section 3.1 then takes the form $T_{t+10} = (1 - \sigma_{forc}^{dec})T_t + \Gamma$, for some $\Gamma \in \mathbb{R}$, and downscaling to an annual time step implies

$$\sigma_{forc}^{ann} = 1 - (1 - \sigma_{forc}^{dec})^{\frac{1}{10}} \approx 0.032 \ , \tag{20}$$

where we obtain $\sigma_{forc}^{dec} = C1 * \lambda \approx 0.28$ directly from the original DICE parametrization. Similarly, we obtain the parameter

$$\sigma_{ocean}^{ann} = 1 - (1 - \sigma_{ocean}^{dec})^{\frac{1}{10}} \approx 0.007 \ . \tag{21}$$

The stationarity assumptions underlying the downscaling are not generally satisfied along an optimal trajectory, but seem the most natural, mechanical approach without recalibrating the model. For a one year time step, however, our calibration-based approach in section 3.2 finds values very close to those deriving from the downscaling argument laid out here. The atmosphere ocean difference ΔT_t is independent of the time step, and equations (14) and (15) are all we need for a rescaling of the time step.

The rate of excess carbon removal from the atmosphere $\delta_{M,t}^{dec}(t, M)$ characterizes decadal decay instead of the annual rate entering equation (12). Similar to equations (20) and (21), we obtain

$$\delta_{M,t}(t, M) = 1 - \left[1 - \delta_{M,t}^{dec}(t, M)\right]^{\frac{1}{10}} \tag{22}$$

by rescaling the corresponding decay factor.

We will employ the sophisticated calibration discussed here to compare the original DICE dynamics on a decadal time scale to that of the interpolated model on an annual time step. Therefore, we will change our pick of $g_{A,0} = g_{A,0}^{DICE}/10$ discussed in section 2.1 by a value that reproduces the DICE model's technology dynamics, including what we argue is likely a rounding error in setting up the DICE 2007 model. We calculate this value as $g_{A,0} = -\frac{\delta_A \log(1-10*(1-\kappa)g_{A,0}^{DICE})}{(1-\kappa)(1-\exp(-10\delta_A))} = 1.46\%$.[11]

4 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.

4.1 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 the value functions V. In the case of uncertainty, the value function generally relies on additional states summarized in the vector $\mathbf{\Phi}_t$, capturing uncertain (formerly) exogenous states, Bayesian information, or persistent components of stochastic shocks (see section 2.3). We then write the value function as $V(K_t, M_t, T_t, \mathbf{\Phi}_t, t)$.

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 η. We denote the rate of pure time preference, also called utility discount rate, by δ_u. The

[11] The DICE model uses $A_{t+10}^{TFP} = \frac{A_t^{TFP}}{1-10g_{A,t}^{TFP}}$ in total factor productivity units. Using labor augmenting technological progress and picking $t = 0$ we obtain $A_{10} = \frac{A_t}{(1-10g_{A,t}^{DICE}(1-\kappa))^{\frac{1}{1-\kappa}}}$, where $g_{A,t}^{DICE} = \frac{g_{A,t}^{TFP}}{1-\kappa}$ is the original DICE model's initial growth rate when converting into labor augmenting technological growth. Setting this equation for A_{10} equal to our continuous time expression in equation (4) delivers the cited formula for $g_{A,0}$.

Bellman equation for standard (entangled) preferences is

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

where E takes expectations over uncertainty in the equations of motion governing period $t + \Delta t$ states. 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 (10), (12), and (13), or their modifications including uncertainty discussed in section 2.3, 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

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

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

$$V^*(k_t, M_t, T_t, \mathbf{\Phi}_{t+\Delta t}, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta}\Delta t + \\ \frac{\beta_{t,\Delta t}}{1-\eta} \; \mathrm{E}\left[V^*(k_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, \mathbf{\Phi}_{t+\Delta t}, t+\Delta t)\right]. \tag{23}$$

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

$$\beta_{t,\Delta t} = \exp\left[(-\delta_u + g_{A,t}(1-\eta) + g_{L,t})\Delta t\right] \; . \tag{24}$$

Its time dependence arises because of the non-constant growth rates in DICE. The factor determines the contraction of the Bellman equation.[12] We can either use the

[12]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.

exact, time-step dependent formulas for the growth rates

$$g_{L,t} = [\ln(L_{t+\Delta t}) - \ln(L_t)]/\Delta t \quad \text{and} \tag{1*}$$

$$g_{A,t} = [\ln(A_{t+\Delta t}) - \ln(A_t)]/\Delta t \ , \tag{3*}$$

or their continuous time approximations (1′) and (3′) at the expense of a small error.[13]

We derive the normalized Bellman equation (23) in Appendix A for the general case of Epstein-Zin preferences. A co-benefit of solving for the value function, 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.[14]

4.2 Solving the model

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

$$\begin{pmatrix} 1 & \dfrac{k_t^\kappa}{1+b_1 T_t^{b_2}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} c \\ \Lambda \end{pmatrix} \leq \begin{pmatrix} \dfrac{k_t^\kappa}{1+b_1 T_t^{b_2}} \\ \Psi_t \end{pmatrix} \quad \text{and } c, \Lambda \geq 0\ . \tag{25}$$

Apart from the physical states, we have to approximate the value function over the 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

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

[13] 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. 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.

[14] The DICE-2007 model 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.

We refer to τ as artificial time. We then generate the grid on the time axis using e.g. Chebychev nodes on $[0, 1)$.[15] A larger choice of the numerical parameter ζ moves the nodes closer to the early (real-time) periods.[16] After generating the nodes, we transform them back to real time using the inverse transformation $t = -\frac{\ln[1-\tau]}{\zeta}$.

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 (23) coincide. We obtain a solution to equation (23) by function iteration, approximating the value function V by a set of basis functions. We recommend Chebychev polynomials. A frequently used alternative are cubic splines, and section 5.3 compares the two approaches, showing that the Chebychev basis requires a lower dimensional bases for the same numerical approximation quality. In Matlab, the compecon toolbox 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 suggested 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.[17]
4. Start with an arbitrary guess for the value function or the coefficients characterizing its approximation.

2) The function iteration:

[15] Chebychev nodes do not lie on the boundary of the interval. In contrast, splines usually take the boundary of the support interval as an evaluation node, and we have to choose an upper bound strictly smaller but close to unity.

[16] 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 ζ is necessary to moderate the clustering at early times and spread the support to a sufficient range in real time. The parameter ζ 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.

[17] 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.

1. Solve the right hand side optimization problem of the Bellman equation (23) 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 functions, and changes in the interval bounds still affect the 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 or cubic splines, a low order of the basis generally results in major wave patterns in the control rules. These patterns are generally of numeric origin. Some modifications of the basic function iteration algorithm described above can be used to speed up convergence. The most important acceleration derives from an algorithm known as Howard's method or modified policy iteration: after every value function iteration, we iterate the Bellman equation several times 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.[18]

[18]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 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). See Cai, Judd & Lontzek (2012*a*) for an example of solving the DICE model over a finite time horizon.

4.3 Epstein-Zin preferences

The Bellman equation in section 4.1 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 (n.d.) argue in 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 η. In an intergenerational interpretation, the parameter η can also be interpreted as the parameter of intergenerational consumption smoothing. Appendix A derives the Bellman equation disentangling risk aversion from intertemporal substitution

$$V^*(k_t, M_t, T_t, \mathbf{\Phi}_{t+\Delta t}, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta}\Delta t + \frac{\beta_{t,\Delta_t}}{1-\eta}\left(\mathrm{E}\left[(1-\eta)V^*(k_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, \mathbf{\Phi}_{t+\Delta t}, t+\Delta t)\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}. \tag{26}$$

We solve the generalized Bellman equation (26) the same way as the original Bellman equation (23), and V^* has the same normalization. Equation (26) 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 (23) vanish and we are back in the setting of equation (23). 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), Bansal, 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 η 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 (n.d.) analyzes the effects of preference disentanglement

under damage uncertainty in detail, and Jensen & Traeger (2013*a*) analyze the effects of preference disentanglement in the context of growth uncertainty.

5 Calibration Results

We now present the results of our calibration, the comparison to DICE and the scientific AOGCMs, and numerical insights regarding the type and number of basis functions. First, we present the result of our simple base case calibration discussed in section 3.2. Then, we compare the temperature response in this simplified model, the original DICE 2007 model, and the DICE 2013 model to the temperature response generated by the scientific AOGCMs for the IPCC's emission scenarios. Finally, we present the result of our sophisticated calibration discussed in section 3.3, and compare the numerical quality of the results across different types and dimensionalities of the basis functions.

5.1 The simple base case calibration

Figure 2 shows our simple base case calibration, resulting from the approach discussed in section 3.2. Our main focus in calibrating was matching DICE's mitigation policies, i.e., the abatement rate and the optimal carbon tax. Our calibration matches optimal policies, and the trajectory of the optimal carbon stock well, with slightly too low a carbon stock four centuries from today. Our temperature evolution shows a slight overshooting two to four centuries from today. It is the immediate result of the simple min-quadratic function used for ocean cooling discussed in section 3.2 (see also Figure 8 in the appendix). The social cost of carbon extends beyond the DICE values when full abatement is optimal. The DICE model calculates the optimal carbon tax using the abatement costs. Once reaching full abatement, this optimal carbon tax is the minimal tax the regulator has to impose on carbon to keep full abatement in place. This crossed line slowly falls over time as abatement becomes cheaper (see equation 6). Our dynamic programming solution obtains the social cost of carbon directly from the value function and, therefore, we can trace out the actual social cost of a ton of atmospheric carbon also beyond the point of full abatement. It then represents the value of capturing a ton of carbon from the atmosphere.

5.2 Base case calibration, DICE, and MAGICC

This section compares our simple base case calibration with exogenous, time-dependent rates of atmospheric carbon removal and ocean cooling to the performance of the climate modules in the DICE 2007 and the DICE 2013 models. As benchmark, we use the MAGICC 6.0 model. The Working Group I of the IPCC (2013) uses MAGICC for Projections of Global and Regional Climate Change (chapter 5), and DICE itself is calibrated to an earlier version of the MAGICC model. MAGICC emulates the state

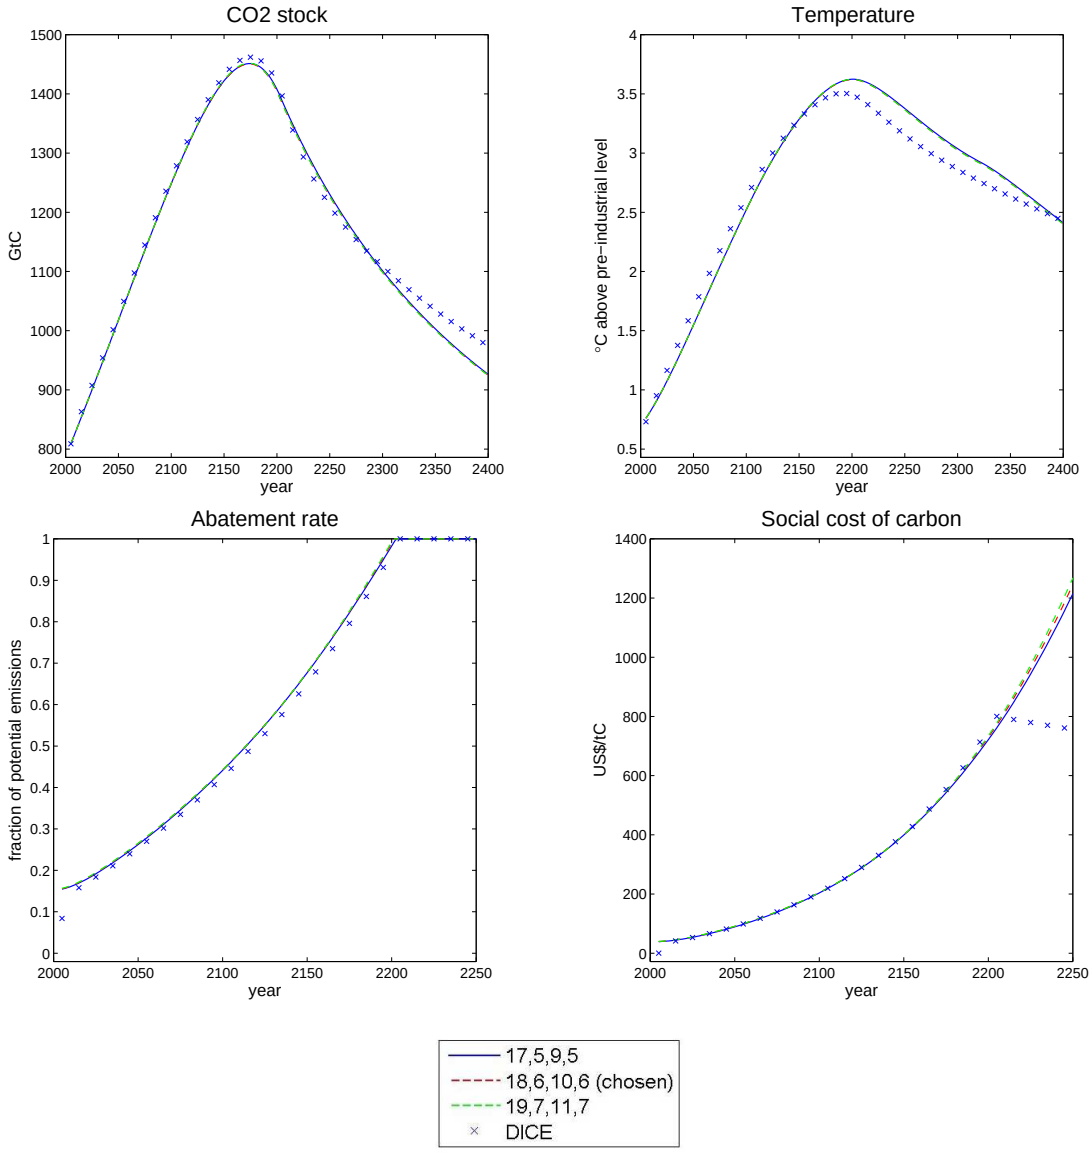

Figure 2 shows the result of our base case calibration to the DICE 2007 model along the deterministically optimal trajectory. We calibrated simple parametric functions to capture a time dependent rate of carbon removal from the atmosphere as well as a simplified temperature delay equation omitting the ocean temperature state. Section 3.2 discusses our simplifications and the calibration approach in detail. Appendix B presents our calibration graphs. Our calibration focuses on reproducing DICE's optimal policy. The different lines reflect different numbers of basis functions (effective capital, carbon stock, temperature, time). They mostly coincide, showing that further increases in the number of basis functions no longer change the result. 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.

of the art climate models. We employ its baseline calibration III (Meinshausen, Raper & Wigley 2011). Figure 4 in Meinshausen et al. (2011) shows that the MAGICC 6.0 emulation replicates well the mean response of 19 Atmosphere-Ocean Global Circulation Model (AOGCM) models employed in the Fourth Assessment Report IPCC (2007) for different policy scenarios.[19] MAGICC is an upwelling diffusion model building on a hemispherically averaged energy-balance equation. It models carbon uptake and (both transient and long-run) warming feedbacks in much more detail than DICE's simple 3 box carbon cycle and temperature delay equations.

DICE and our integrated assessment model only endogenize fossil fuel bases CO_2 emissions. CO_2 emissions from land use change and forestry follow the exogenous trajectory described by equation (7), and all other greenhouse gases are summarized in the exogenous radiative forcing term characterized by equation (8). In contrast, the MAGICC model, like the big, scientific climate change models, explicitly models the dynamics of a large set of greenhouse gases. The IPCC's emission scenarios vary CO_2 emissions as well as the emission levels of other greenhouse gases. Indeed, stricter mitigation policies would not only regulate industrial CO_2 emissions more strictly, but also other greenhouse gas emissions. Our comparison presented here stacks the cards against DICE and our model: we show how well our model represents the different SRES and RCP scenarios if we only adjust endogenous CO_2 emissions. We thereby answer the question how well the simplified integrated assessment models, endogenizing only the main source of greenhouse gas emissions, replicate the responses to more comprehensive policy approaches regulating greenhouse gases in general. Appendix C presents the complementary approach: Figures 11 and 12 compare the temperature response from DICE and our model to MAGICC when feeding the radiative forcing implied by all greenhouse gases in the SRES and RCP scenarios. In this comparison, DICE 2007, DICE 2013, and our base case model perform significantly better in reproducing MAGICC's temperature response as compared to the case presented here. Moreover, in the majority of cases our model performs better than DICE 2007 or DICE 2013.

Figure 3 compares our base case calibration from section 3.2 to the emission response of the MAGICC 6.0 model. We selected 3 common SRES scenarios used in the IPCC's assessment reports up to Assessment Report 4 (left), and 3 RCP scenarios used in the newest Assessment Report 5 (right). The new RCP scenarios have an extended time horizon as compared the 100 year time horizon of the SRES scenarios. Figure 12 in Appendix C presents the comparison for the remaining scenarios.[20] All

[19]Fourteen modeling groups had submitted data for 23 AOGCMs and MAGICC 6.0 is calibrated to the 19 of these AOGCMs for which sufficient data was available to carry out the calibration. Calibration III uses the widest set of calibration runs. See section 4.3.2 of Meinshausen et al. (2011) for a discussion of the emulation error using MAGICC 6.0. For individual AOGCMs the the error is small as compared to the disagreement of individual AOGCMs. For the purpose of the current paper, I am interested in the model's ability to emulate the mean response of the AOGCMs, where MAGICC 6.0 performs even better.

[20]See IPCC (2000) for details on the scenarios. They reflect a wide range of emission forecasts

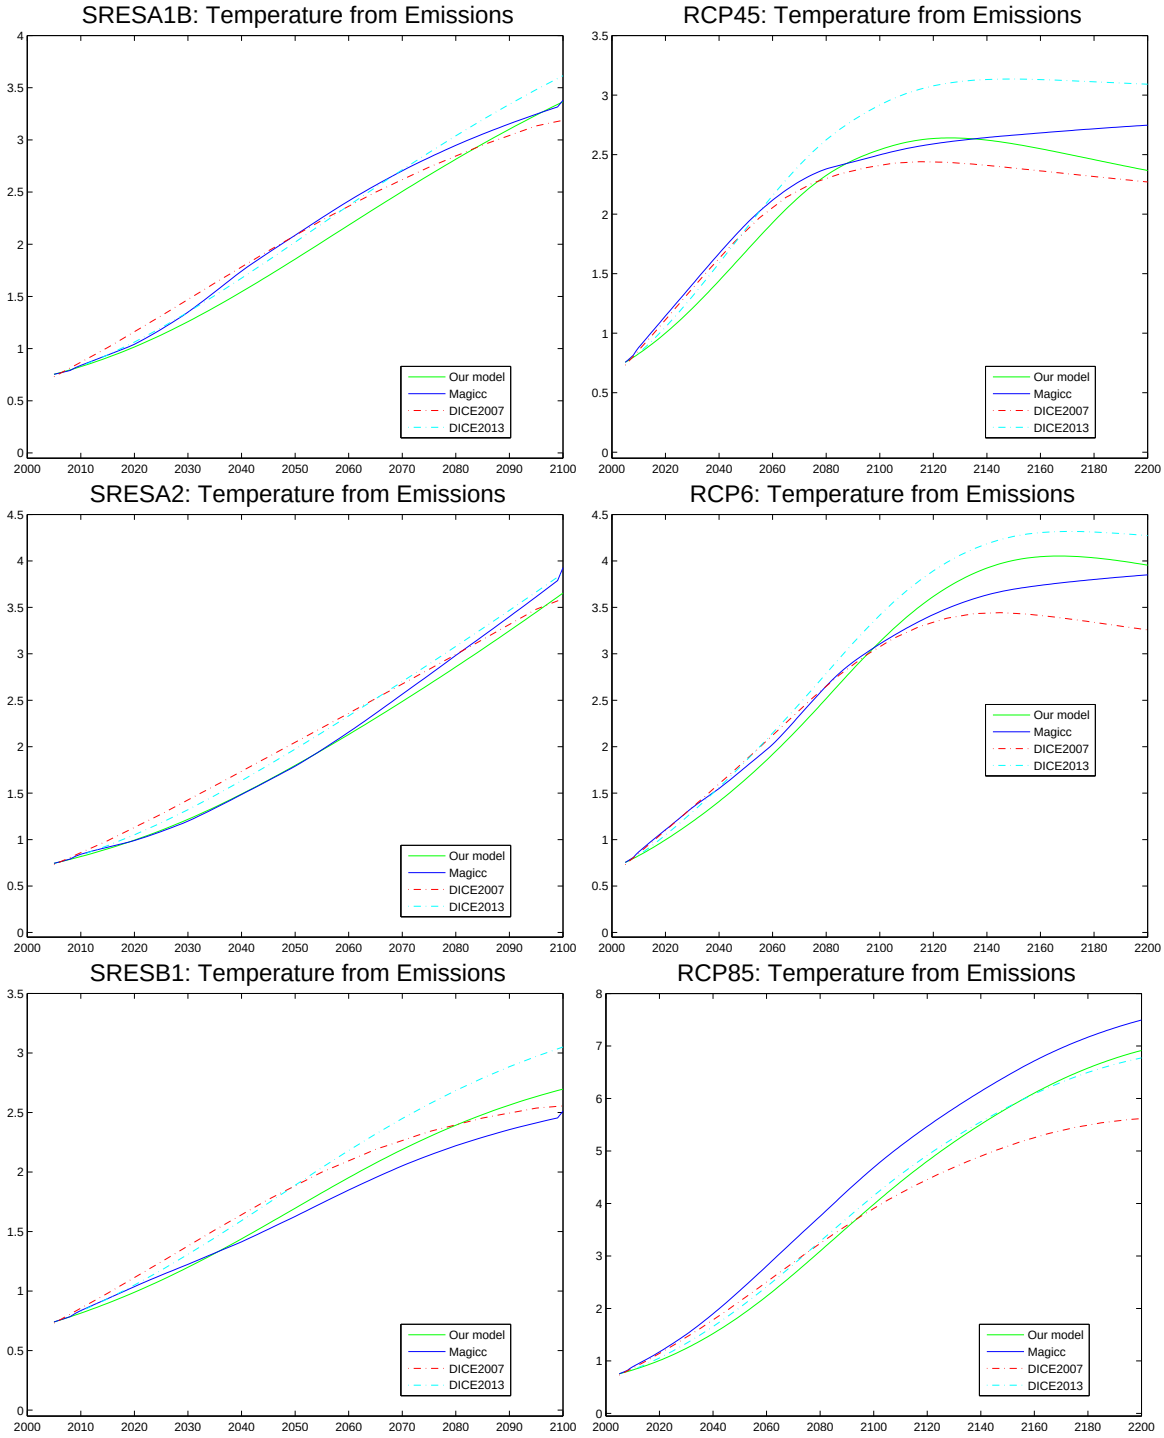

Figure 3 compares the temperature response of our simple calibration, of DICE 2007, and of DICE 2013 to MAGICC's temperature response. We show selected emission scenarios used in the 4^{th} (SRES, left) and 5^{th} (RCP, right) assessment reports of the IPCC. Appendix C presents additional results.

integrated assessment models resemble MAGICC's temperature response to emission reasonably well. While there clearly is room for improvement, we find that our much simplified model is no worse in reproducing temperature responses for a large set of emission scenarios than either of the DICE 2007 or the DICE 2013 models. We observe the strongest deviation from MAGICC in Figure 3 for the RCP 8.5 scenario. The scenario exceeds 7°C temperature increase during the next century and is far from an optimal emission trajectory that our model is designed to generate. The largest deviation to both the original DICE and our base case model is generated by the RCP 3 scenario depicted in Figure 12 in Appendix C. This RCP 3 scenario is MAGICC's version of the IPCC's RCP 2.6 scenario and assumes extreme, very quick emission reductions. By the end of the century fossil fuel based emissions are negative (net carbon capture). Neither DICE nor our model are designed to simulate such an extreme emission reduction scenario.

5.3 Sophisticated calibration, Chebychev polynomials, splines

In this section, we briefly present optimal abatement in the more sophisticated interpolation of the DICE model presented in section 3.3. As we discussed there, the climate response necessarily differs for different non-stationary emission trajectories between decadal and annual models. In addition, despite our careful interpolation of DICE's exogenous processes, the dynamics of our, endogenous, downscaled, economic equations of motions will not exactly match those of the decadal model. In addition, our annual time step allows the policy maker to set a slightly better targeted policy with increased efficiency.

The solid blue lines in Figure 4 compare the optimal policies between the original DICE 2007 model and our sophisticated interpolation in an annual model. These solid blue lines are partially covered by a dashed yellow line showing that further increasing the order of the basis does not change the numerical results. The left panels present the results using a Chebychev basis, and the right panels present the solutions using a cubic spline basis. We discuss the numerical insights from the graphs further below and start discussing the relation between the crossed line (markers) representing the original decadal DICE model and our solid lines for an annual time step. The upper two panels in Figure 4 show that the optimal abatement rate and the optimal emission levels in the present coincide for both models. Later, both the abatement rate and the absolute emission levels fall slightly below the DICE values, returning to the DICE values during the next century. The main reason why both abatement rate and emission levels can fall below their DICE counterparts is that the capital stock grows

that represent different assumptions on population growth, technological development and cross-regional spill over, and economic development and convergence. The new RCP scenario are labeled by the radiative forcing they produce by the end of the century and are described in Moss, Babiker, Brinkman, Calvo, Carter, Edmonds, Elgizouli, Emori, Erda, Hibbard, Jones, Kainuma, Kelleher, Lamarque, Manning, Matthews, Meehl, Meyer, Mitchell, Nakicenovic, O'Neill, Pichs, Riahi, Rose, Runci, Stouffer, van Vuuren, Weyant, Wilbanks, van Ypersele & Zurek (2007).

at a slightly slower rate, a consequence mostly of the capital decay change along the non-stationary path in the annual time step.[21] The main message from comparing the more sophisticated interpolation and the original DICE model is that, first, a sophisticated interpolation of DICE in combination with a more mechanical downscaling of the time step will not necessarily imply the same policy and temperature dynamics as the original model. Second, these deviations, however, are not particularly large. We note that paralleling research by Cai et al. (2012*a*), Lemoine & Traeger (2014), Cai, Judd & Lontzek (2012*c*), and Alex L. Marten (2013) also discusses the consequences of the time step on integrated assessment. For the purpose of the current paper we take the following stand. The current paper's goal is to promote a model that helps analyzing the policy implications of uncertainty in climatic change. Our simple calibration discussed in sections 3.2, 5.1, and 5.2 perfectly resembles DICE's optimal policy in a deterministic world, and captures reasonably well the temperature response to a wide set of emission scenarios. We consider the simple calibration a perfectly suitable tool for our suggested purpose: analysis building on the model will capture the response of optimal policy to uncertainty. Such analysis derives its insights relative to the original DICE model's deterministic policy. The convenience of having a model that reproduces the DICE model's deterministic policy and that is easy to implement outweighs the costs of not taking a more sophisticated stand on optimal policy changes from downscaling the time step in a deterministic world.

Figure 4 also discusses numerical aspects of the solutions. The legend entries state the numbers of basis functions in each dimension. The first number corresponds to the effective capital state, the second number to the carbon stock, the third number to the temperature state, and the fourth number to the time state. We use a tensor basis, i.e., the basis of our full space consists of all possible combinations of basis functions in each dimension. Thus, the number of basis functions approximating the value functions is the product of the different legend entries, which is also the number of nodes on our supporting grid. First and most importantly, we show that our results are independent of our choice of the basis type. Once we have picked a sufficiently large number of basis functions, the solid lines on the left using a Chebychev basis closely match the solid lines on the right using a spline basis. Second, we show that using too low a number of basis functions can imply significant numerical mistakes in the solution (dotted, and green and red dashed lines). Third, the figure shows that Chebychev polynomials yield the same quality of the solution with a significantly lower number of basis functions than splines. We suggest to employ the number of basis functions stated next to the solid blue line. Further increasing the number of basis functions (dashed yellow line) changes the results so slightly that the yellow

[21] A small difference also results from the fact that the EXCEL spreadsheet version of DICE that we use for our comparison differs slightly from the GAMS version that we implemented in Matlab. We realized the EXCEL version of DICE subtracts abatement costs in a way that is independent of climate damages, whereas abatement costs in the GAMS and, thus, our version of DICE scale with production net of climate damages. The effect of this difference is very small (proportional to damages times abatement cost, both measured as percent of world output).

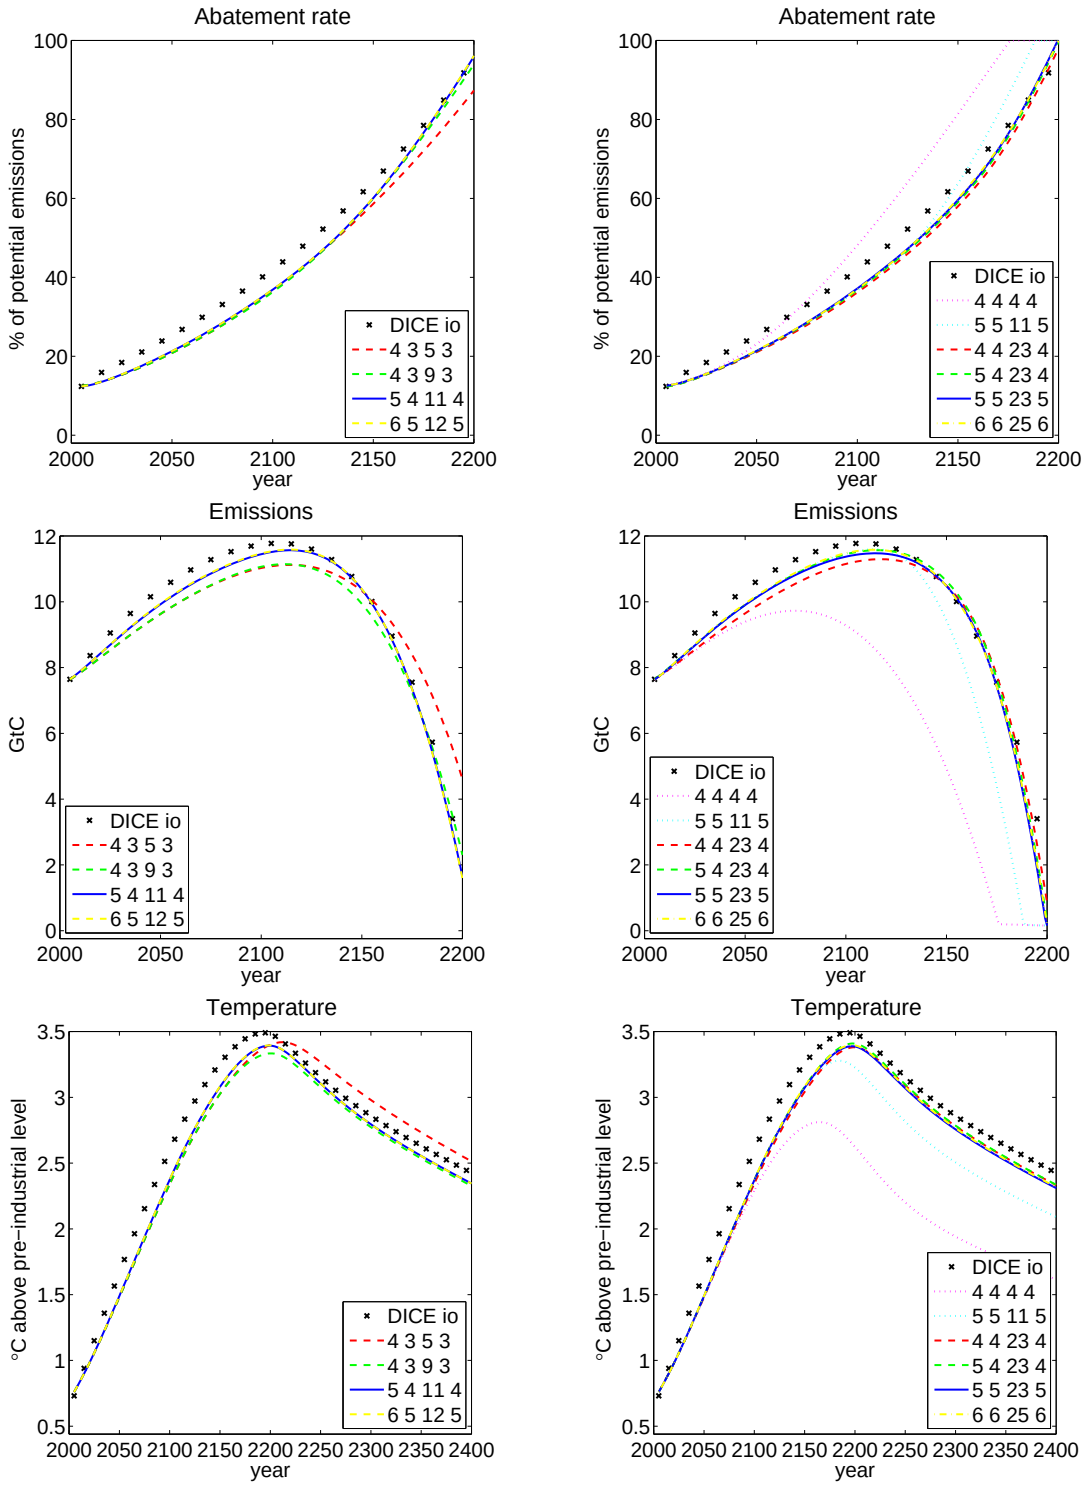

Figure 4 presents the optimal emission policy and temperature evolution in our "sophisticated" DICE 2007 interpolation (section 3.3) for an annual time step. The graphs on the left use a Chebychev basis, whereas the graphs on the right use a spline basis. The legend entries state the numbers of basis functions in each dimension: effective capital, carbon stock, temperature, time. The yellow dashed lines mostly lie on top of the solid blue lines.

dashed line lies essentially on top of the solid blue line, making it appear a single dashed line. In the case of Chebychev polynomials we find that $5 \times 4 \times 11 \times 4 = 880$ basis functions already yield a very good numerical approximation to the value function, while in the case of splines we need $5 \times 5 \times 23 \times 5 = 2875$ basis functions to reach a similarly good solution. Using collocation, the number of evaluation nodes equals the number of basis functions, and processor time is approximately proportional to the number of evaluation nodes.

The state space underlying Figure 4 uses a the support $[3, 6]$ for effective capital, the support $[550, 2000]$ in Gt carbon for the carbon stock, the support $[0, 5]$ for the temperature increase in °C above temperature in 1900, and the support $[0, \infty)$ for Chebychev polynomials and $[0, 576]$ in the case of splines for real time in years into the future.[22] The fourth finding from our figure is that, on this support, the time dimension needs by far the highest dimensional basis to obtain a good approximation. A reason for the finding is that the DICE 2007 model has a significant amount of exogenous processes that drive a non-stationary dynamics.[23] The number of basis functions needed for a precise value function approximation depends on the size of the support intervals. Uncertainty can carry trajectories out of these supports. E.g., if uncertainty about climate sensitivity implies significant probabilities that we reach higher temperature levels, we have to extend the support interval, and usually as well the order of the basis in the temperature dimension. Our base case calibration shown in Figure 2 prepared the model for growth uncertainty, in which case we need a much larger effective capital state space because technological growth follows e.g. a random walk. There, we used an effective capital state space of $[0.5, 15]$ and ended up using 18 rather than 5 capital nodes until the solution no longer depended on further increasing the basis. In particular, this finding illustrates the convenience of our renormalization of the model and Bellman equation in per effective labor units. Our trajectories underlying Figure 4 stay well inside the chosen support interval of $[3, 6]$ over the full time horizon from the present to infinity. Without normalizing

[22] These time intervals correspond to $[0, 1)$ and $[0, 0.99999]$ in artificial time, respectively. Whereas splines place a node on the bounds of the support interval, the Chebychev nodes lie stricly inside of the support interval, thus, allowing us to use infinity as the upper bound. Also in the case of the Chebychev basis our actual time nodes do not exceed the 600 year time horizon that our sophisticated interpolation relies on for interpolating DICE's original emission and temperature dynamics. However, we show that further increasing the node numbers in the time dimension has a negligible impact on the model results. Increasing the node number also pushes the highest Chebychev nodes further out into the future. Note that placing the highest time node at year 576 does not imply that the policy maker's planning horizon ends there. We use a smooth interpolation in the time dimension, not cutting off the planning horizon at any given point.

[23] Note that our logarithmic transformation clusters more nodes on in the close future as opposed the distant future. We found that this clustering is useful because the DICE model is particularly non-stationary in the close future. We picked the current transformation after testing a small set of different time transformations, including a logarithmic variation that reduces the clustering. However, it is more than likely that other time transformations can be found that reduce the number of time nodes needed to obtain the same quality in the approximations. Note that for some transformations the model becomes less stable than under the logarithmic transformation.

capital to effective labor units we would need a much larger state space for capital to cover at least a reasonably long time horizon, even without growth uncertainty.

6 Conclusions

The current paper develops a low dimensional model for analyzing the policy implications of uncertainty in climatic change. In contrast to Monte-Carlo studies, linearizations, or discretized models, our dynamic programming approach solves the non-linear problem of finding the optimal policy in the face of uncertainty. The present model relates closely to the wide-spread DICE model, but uses an annual, or more generally flexible, time step. We interpolate the complex set of exogenous difference equations in the recent DICE model by their continuous time solutions, making this integrated assessment model even more accessible to introspection by a large audience of environmental economists.

Our model extends the small set of existing state of the art implementations of stochastic dynamic programming integrated assessment models. Our main contribution is to reduce the number of states needed to represent the climate side of the current DICE model without sacrificing its benchmark in capturing the interaction between emissions and temperature increase. Our reduction of the state space is crucial to permit additional state variables needed to capture uncertainty. We cut the number of state variables almost to half with respect to DICE 2007 and DICE 2013. We compare the performance of our simple model, and of the original DICE 2007 and DICE 2013 models, to the temperature response of MAGICC 6.0. MAGICC emulates the Atmosphere-Ocean General Circulation Models used in the IPCC's assessment reports. Our scenario comparison uses the IPCC's SRES and RCP emission scenarios and finds that our simplified model performs similarly well as the original DICE model. In general, the accuracy of these integrated assessment models in emulating the scientific AOGCM models can be improved.

Our paper is the first to carefully analyze the numerical quality of a stochastic dynamic programming implementation of an integrated assessment model, comparing results for different types of basis functions and determining the number of basis functions needed for an accurate solution. In our DICE implementation, the Chebychev basis reaches numeric precision with significantly less basis functions as compared to cubic splines. In addition, we introduce a new, efficient normalization of the Bellman equation and introduce Epstein-Zin preferences. These preferences allow the model to distinguish effects of risk and risk aversion from those deriving from intertemporal consumption smoothing. Our flexible time step thereby prepares the model to use the recent consumption smoothing and risk aversion estimates from the finance literature.

References

Alex L. Marten, S. C. N. (2013), 'Temporal resolution and dice', *Nature Climate Change* **3**(6), 526.

Bansal, 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.

Cai, Y., Judd, K. L. & Lontzek, T. S. (2012*a*), Continuous-time methods for integrated assessment models, NBER Working Papers 18365, National Bureau of Economic Research, Inc.

Cai, Y., Judd, K. L. & Lontzek, T. S. (2012*b*), Dsice: A dynamic stochastic integrated model of climate and economy, Working Paper 12-02, RDCEP.

Cai, Y., Judd, K. L. & Lontzek, T. S. (2012*c*), 'Open science is necessary', *Nature Climate Change* **2**, 299.

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

Crost, B. & Traeger, C. (n.d.), 'Risk and aversion in assessing climate change'.

Crost, B. & Traeger, C. P. (2013), 'Optimal climate policy: Uncertainty versus Monte-Carlo', *Economic Letters* **120**, 552–558.

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.

Fischer, C. & Springborn, M. (2011), 'Emissions targets and the real business cycle: Intensity targets versus caps or taxes', *Journal of Environmental Economics and Management* **62**(3), 352 – 366.

Heutel, G. (2011), 'How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks', *University of North Carolina, Greensboro, Department of Economics Working Paper Series* **11**(08).

Hoel, M. & Karp, L. (2001), 'Taxes and quotas for a stock pollutant with multiplicative uncertainty', *Journal of Public Economics* **82**, 91–114.

Hoel, M. & Karp, L. (2002), 'Taxes versus quotas for a stock pollutant', *Resource and Energy Economics* **24**, 367–384.

Interagency Working Group on Social Cost of Carbon, U. S. G. (2010), 'Technical support document: Social cost of carbon for regulatory impact analysis under executive order 12866', *Department of Energy* .

Interagency Working Group on Social Cost of Carbon, U. S. G. (2013), 'Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866', *Department of Energy* .

IPCC (2000), *Emissions Scenarios*, Cambridge University Press, Cambridge, UK.

IPCC (2007), *Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007*, Cambridge University Press, Cambridge, UK.

IPCC (2013), *Climate Change 2013: The Physical Science Basis*, Intergovernmental Panel on Climate Change, Geneva. Preliminary Draft.

Jensen, S. & Traeger, C. (2013*a*), 'Growth uncertainty in the integrated assessment of climate change', *E^3 Working paper WP-065* .

Jensen, S. & Traeger, C. P. (2013*b*), Optimally climate sensitive policy under uncertainty and learning, Working paper, University of California, Berkeley.

Karp, L. & Zhang, J. (2006), 'Regulation with anticipated learning about environmental damages', *Journal of Environmental Economics and Management* **51**, 259–279.

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. (2005), 'Price and quantity regulation in general equilibrium', *Journal of Economic Theory* **125**(1), 36 – 60.

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.

Kelly, D. L. & Tan, Z. (2013), 'Learning and climate feedbacks: Optimal climate insurance and fat tails', *University of Miami Working paper* .

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

Lemoine, D. & Traeger, C. (2014), 'Watch your step: Optimal policy in a tipping climate', *American Economic Journal: Economic Policy* **6**, 1–31.

Meinshausen, M., Raper, S. & Wigley, T. (2011), 'Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, magicc6 - part 1: Model description and calibration', *Atmospheric Chemistry and Physics* **11**, 1417–1456.

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

Moss, R., Babiker, M., Brinkman, S., Calvo, E., Carter, T., Edmonds, J., Elgizouli, I., Emori, S., Erda, L., Hibbard, K., Jones, R., Kainuma, M., Kelleher, J., Lamarque, J. F., Manning, M., Matthews, B., Meehl, J., Meyer, L., Mitchell, J., Nakicenovic, N., O'Neill, B., Pichs, R., Riahi, K., Rose, S., Runci, P., Stouffer, R., van Vuuren, D., Weyant, J., Wilbanks, T., van Ypersele, J. P. & Zurek, M. (2007), *Towards New Scenarios for Analysis of Emissions, Climate Change, Impacts, and Response Strategies.*, Intergovernmental Panel on Climate Change, Geneva.

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.

Nordhaus, W. & Boyer, J. (2000), *Warming the World*, MIT Press, Cambridge, MA.

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

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

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

$$V(k_t A_t L_t, M_t, T_T, \mathbf{\Phi}_t, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} A_t^{1-\eta} L_t \Delta t + \frac{\exp[-\delta_u \Delta t]}{1-\eta} \exp[(1-\eta) g_{A,t} \Delta t] \times$$

$$A_t^{1-\eta} \exp[g_{L,t}\Delta t] L_t \left(\mathrm{E}\left[(1-\eta)\frac{V(K_{t+\Delta t}, M_{t+\Delta t}, T_{t+\Delta t}, \mathbf{\Phi}_{t+\Delta t}, t+\Delta t)}{A_{t+\Delta t}^{1-\eta} L_{t+\Delta t}}\right]^{\frac{1-\mathrm{RRA}}{1-\eta}}\right)^{\frac{1-\eta}{1-\mathrm{RRA}}}.$$

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 divide both sides of the equation by $A_t^{1-\eta} L_t$ and define $V^*(k_\tau, M_\tau, T_\tau, \mathbf{\Phi}_\tau, \tau) = \frac{V(K_t, M_t, T_t, \mathbf{\Phi}_t, t)|_{K_t = k_t A_t L_t}}{A_t^{1-\eta} L_t}$ yielding

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

and, thus, equation (26) and its special case equation (23).

B Calibration

The section summarizes the numerical part of the calibration discussed in section 3.2 whose result we presented in section 5.1. Table 1 summarizes the model parameters. We calibrate our model so that the optimal time paths of CO_2 concentration, temperature, abatement rate, and optimal carbon tax closely resemble those predicted by the DICE-2007 model.[24] Figures 5 and 6 show our calibration of equation (16),

[24] 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

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 δ_M^*. Figure 7 shows the calibration of the heat capacity and feedback related delay parameter σ_{forc}, and of the parameter σ_{ocean} capturing ocean temperature related feedbacks (equation 13). Any individual parameter change improving the fit in one dimension worsens the fit in at least one of the other variables. Finally, Figure 8 depicts the difference between oceanic and atmospheric temperatures in DICE-2007, and compares it to our simple quadratic approximation stated in equation (17).

We solve the Bellman equation (23) by help of the function iteration algorithm described in section 4.2. We approximate the value function V^* by Chebyshev polynomials. We update the basis coefficients by collocation at the Chebychev 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 2 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 10 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.

C Further comparisons to the MAGICC model

The present part of the appendix extends the comparison of our base case calibration and of the DICE 2007 and DICE 2013 models to the average of the AOGCM models used in the IPCC as emulated by MAGICC 6.0. We explained in section 5.2 that MAGICC and the SRES and RCP scenarios model a large number of greenhouse gas emissions explicitly, and that the DICE (and our) integrated assessment model only endogenize fossil fuel based CO_2 emissions. Our comparison in section 5.2 asks how well the integrated assessment models perform in approximating more comprehensive climate policy scenarios. Part of the failure in reproducing the temperatures of the SRES and RCP scenarios is that these simplified integrated assessment models have a fixed exogenous forcing proxy for all non-CO_2 greenhouse gas emissions. A more favorable comparison between the integrated assessment models and MAGICC feeds the joint radiative forcing of all greenhouse gases of the SRES and RCP scenarios to

in Nordhaus (2008). Note that the EXCEL model assumes a constant savings rate. We find an almost constant savings rate in our optimizing model, 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. The EXCEL solver is not able to solve for both investment and abatement over the full trajectory. We have also modified the EXCEL DICE model to explicitly optimize investement over the first century and then jointly chose an investment rate for the remaining time horizon. The impact on optimal abatement was very minor and Figure 4 shows the crosses corresponding to this modified DICE optimization ("DICE io" for "investment optimized").

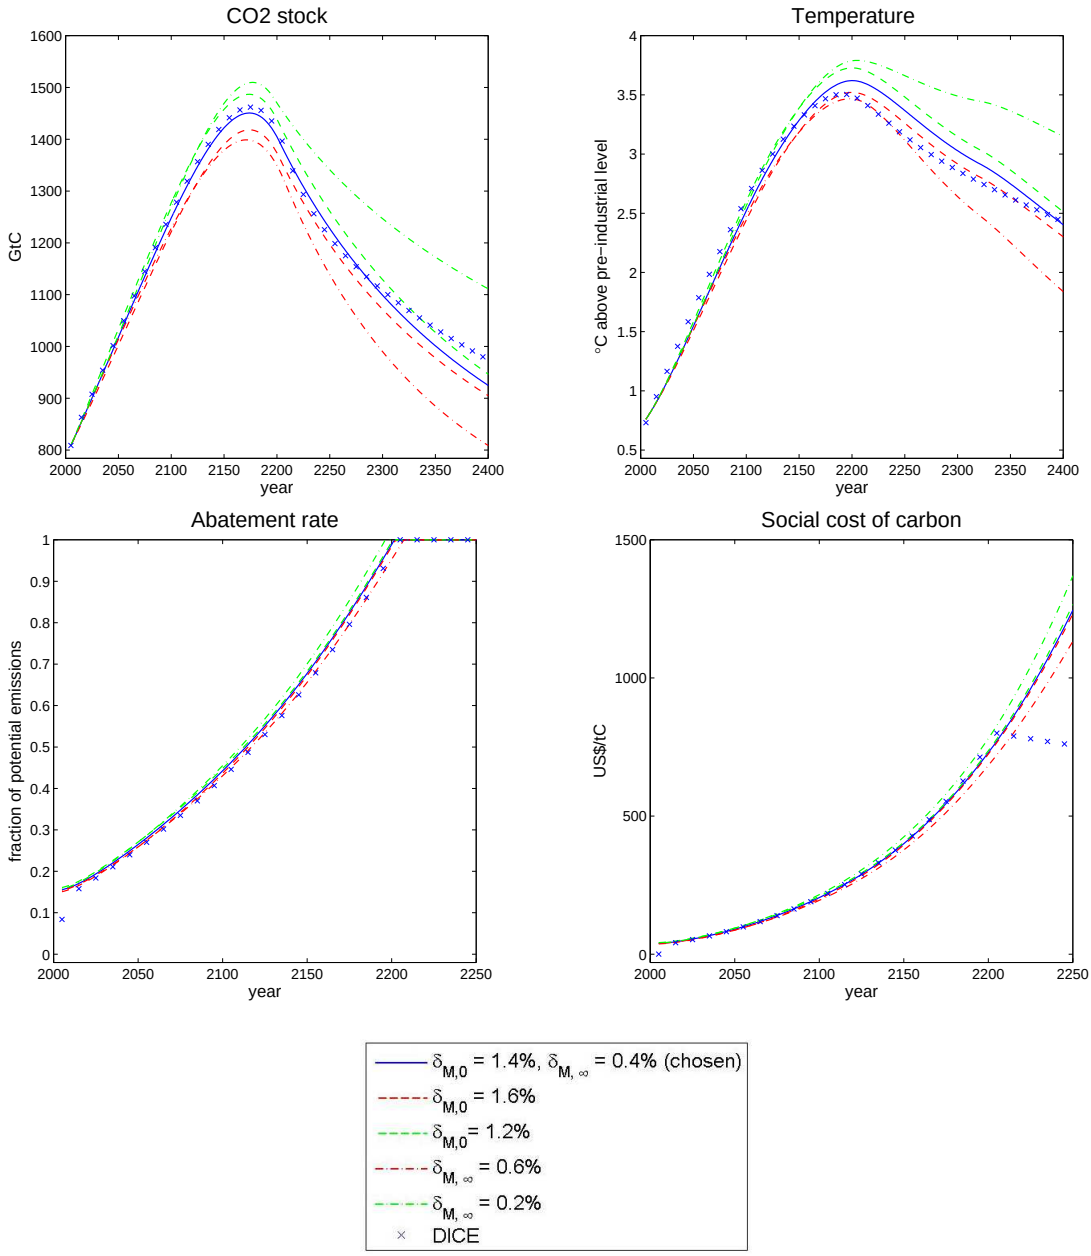

Figure 5 shows the calibration of the rate governing CO_2 removal from the atmosphere. We calibrate the initial rate $\delta_{M,0}$ and the asymptotic rate $\delta_{M,\infty}$. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

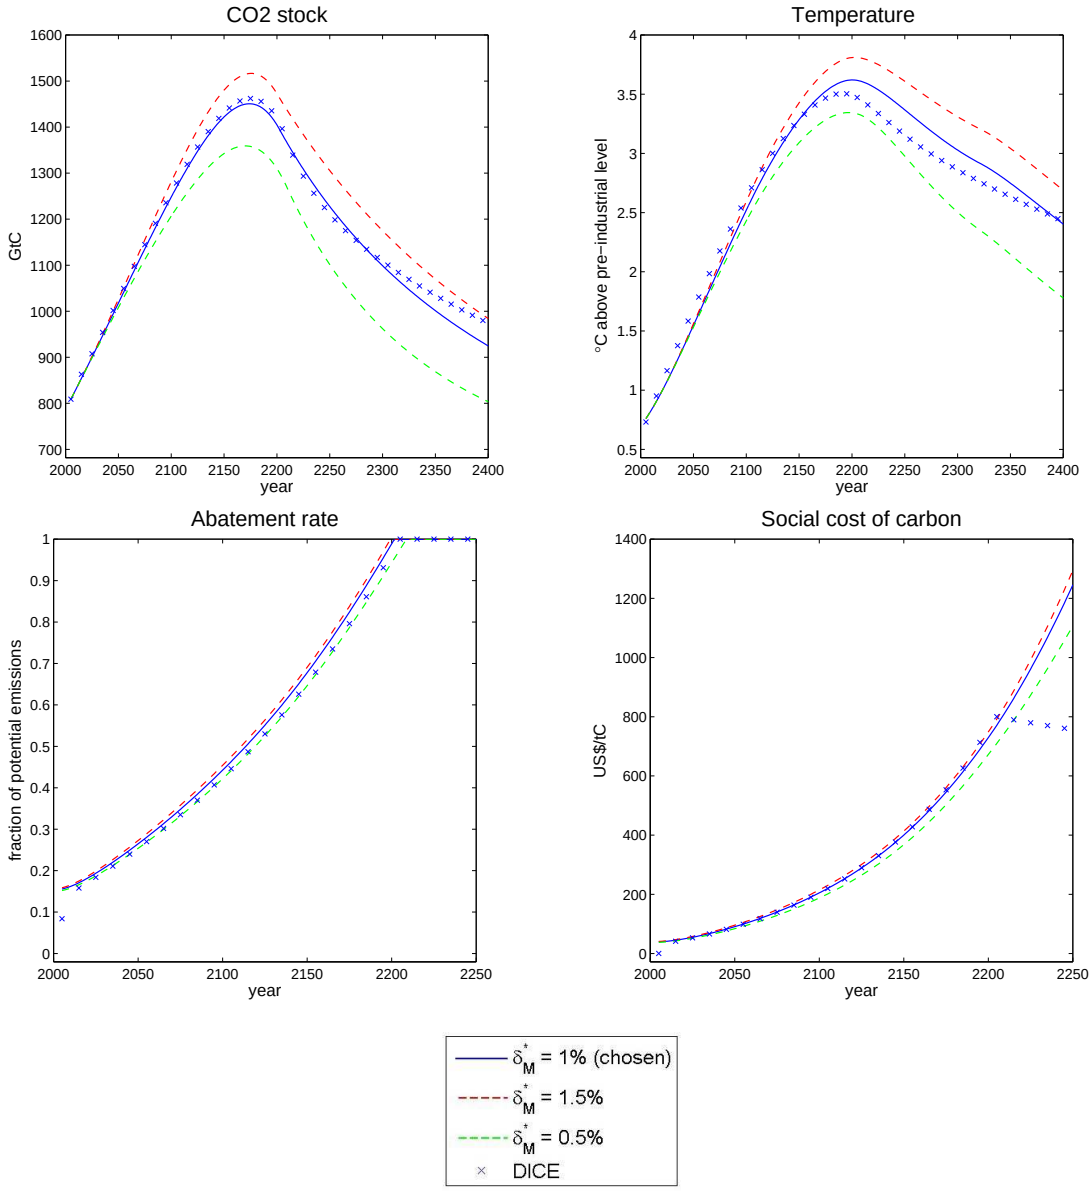

Figure 6 shows the calibration of the rate governing CO_2 removal from the atmosphere. We calibrate the parameter δ_M^* governing the speed of convergence from the initial to the asymptotic rate of CO_2 removal.

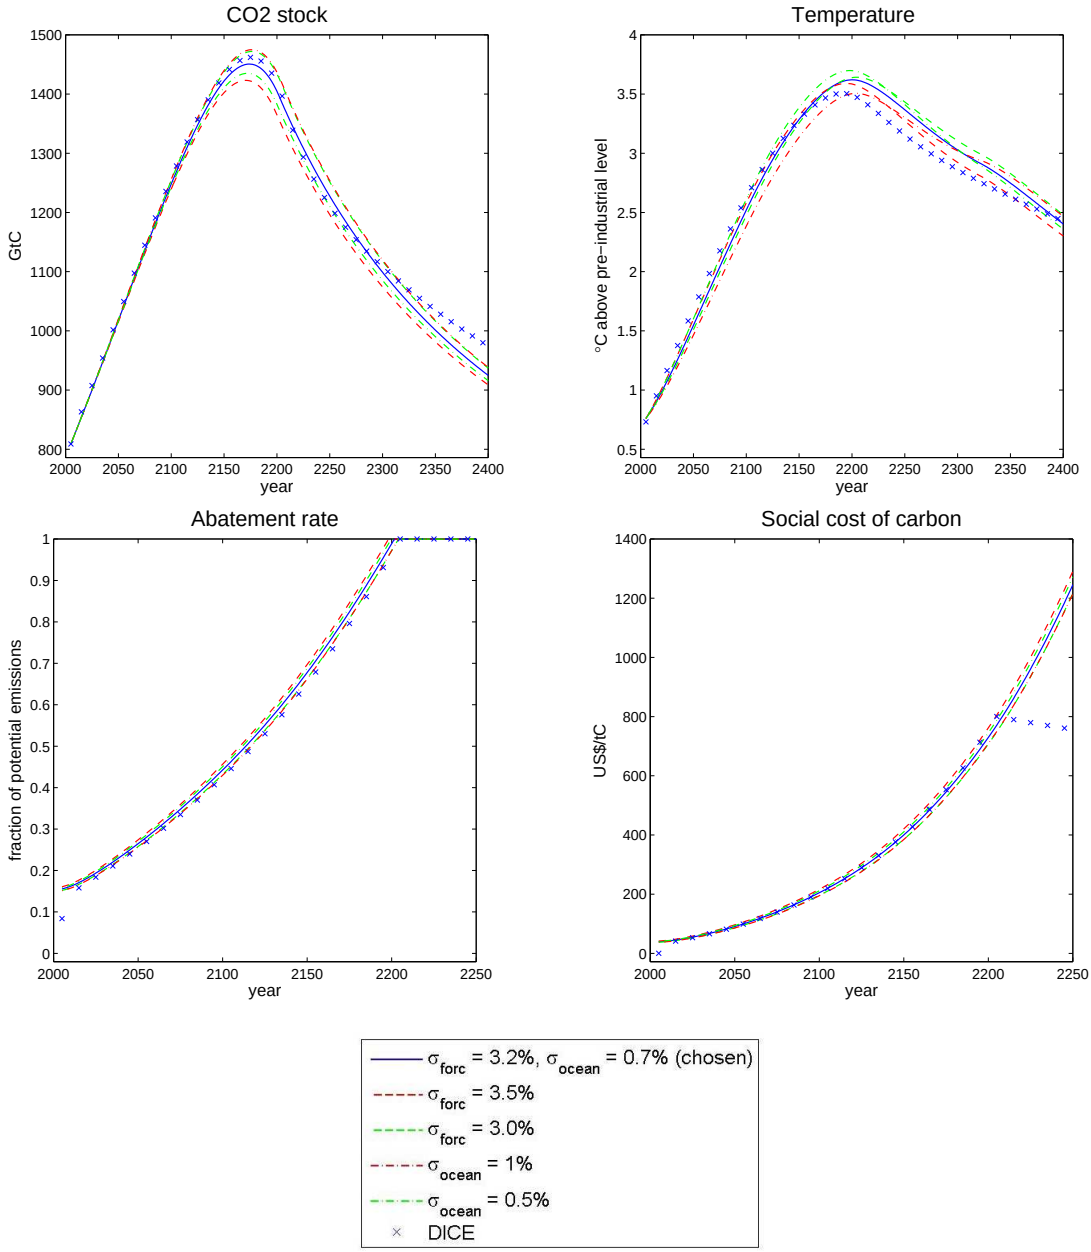

Figure 7 shows the calibration of the warming delay parameter σ_{forc} and the parameter σ_{ocean} connecting atmospheric and oceanic temperatures. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

Table 1 Parameters of the model

Economic Parameters		
η	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
δ_u	1.5%	pure rate of time preference per year
L_0	6514	in millions, population in 2005
L_∞	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
δ_K	10%	depreciation rate of capital per year
κ	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, per year
δ_A	0.1%	rate of decline of productivity growth rate per year
σ_0	0.1342	CO_2 emissions per unit of output in 2005
$g_{\sigma,0}$	-0.73%	initial rate of decarbonization per year
δ_σ	0.3%	rate of decline of the rate of decarbonization per year
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_Ψ^*	-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 CO_2 removal from the atmosphere per year
$\delta_{M,\infty}$	0.4%	asymptotic rate of CO_2 removal from the atmosphere per year
δ_M^*	1%	rate of convergence to asymptotic rate of atmospheric CO_2 removal
B_0	1.1	in GtC, initial CO2 emissions from LUCF
δ_B	1.05%	growth rate of CO2 emisison from LUCF per year
s	3.08	climate sensitivity (equilibrium temperature response to doubling of atmospheric CO_2 concentration w.r.t. preindustrial)
η_{forc}	3.8	forcing of CO_2-doubling
λ	$\eta_{forc}/s \approx 1.23$	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
σ_{forc}	3.2%	warming delay, heat capacity atmosphere, annual
$\tilde{\sigma}_{ocean}$	0.7%	parameter governing oceanic temperature feedback, annual

DICE and our model. Figure 11 presents the results for the same scenarios that Figure 3 compares based on their industrial CO_2 emissions component. As to be expected, all models perform significantly better in reproducing the temperature response of

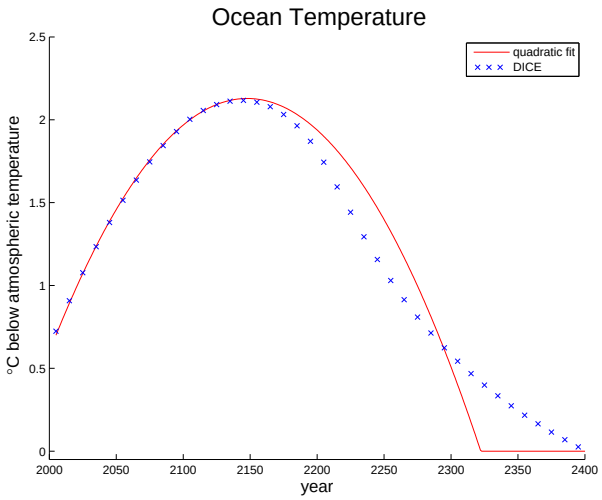

Figure 8 compares our simple min-quadratic approximation of the temperature difference between the atmosphere and the oceans given in equation (17) to the actual difference resulting from the DICE-2007 model. The noticeable difference emerging two centuries into the future causes also the slightly more pronounced difference between our and DICE's atmospheric temperature observed after the year 200 in the earlier calibration plots. Section 5.3 discusses a more sophisticated interpolation using a Chebychev basis and a larger set of interpolation paths (see Figure 9.

Table 2 Location of Collocation Nodes

Node	Effective Capital (k)	Carbon Stock (M)	Transformed Time (τ)	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			

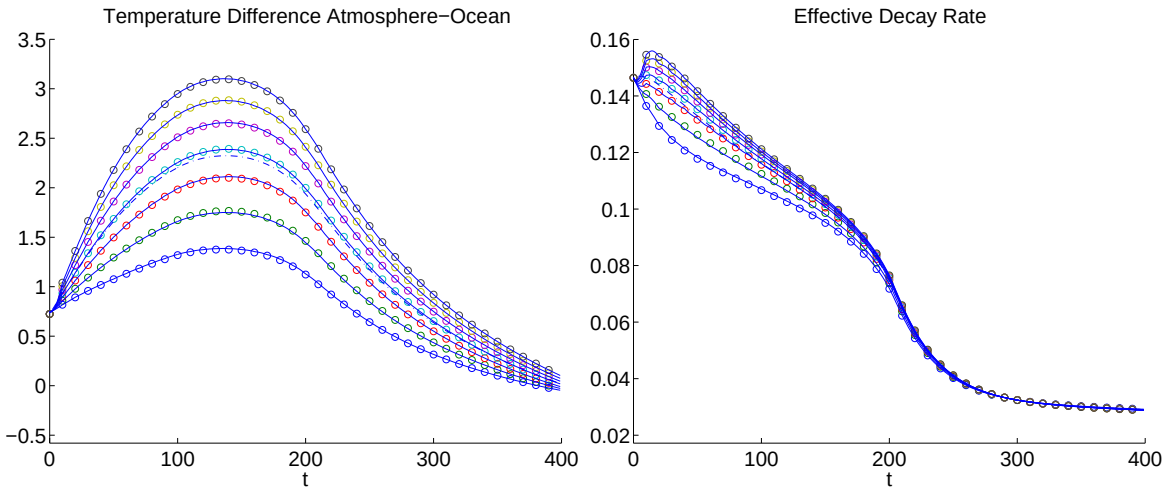

Figure 9 compares our more sophisticated interpolation of the atmosphere-ocean temperature difference and the rate of carbon removal (section 5.3). The circles show the decadal values resulting from the original DICE paths (data). The dashed black line shows the (average) time trend fitted by a 30 node cubic spline. The solid lines recover the original data using the time trend and a linear, period-dependent correction based on the atmospheric carbon content.

MAGICC. In the majority of cases, our simplified model performs slightly better than the original DICE model despite our simplifications of the temperature dynamics.

Figure 12 fills in the same comparison undertaken in Figures 3 and 11 for the remaining SRES and RCP scenarios. The SRES A1F is a more fossil fuel intensive scenario and, like RCP 8.5 in the new scenarios, it is not a candidate for an optimal policy as it significantly exceeds 4°C before the end of the century. RCP 3 is the other extreme. Here CO_2 emissions become negative already in the current century as a consequences of major immediate reductions in combination with carbon capture. Again, it is not a candidate for an optimal policy – without significantly altering abatement cost and damage assumptions and introducing cheap carbon sequestration. The major divergence between our base case calibration and the MAGICC temperature in the emissions based graph on the left hand side is a consequence of the major reduction of non-fossil fuel based emission in MAGICC not accounted for in the integrated assessment models: the graph on the right hand side takes the radiative forcing of all emissions into account and delivers a temperature response closely resembling that of MAGICC.

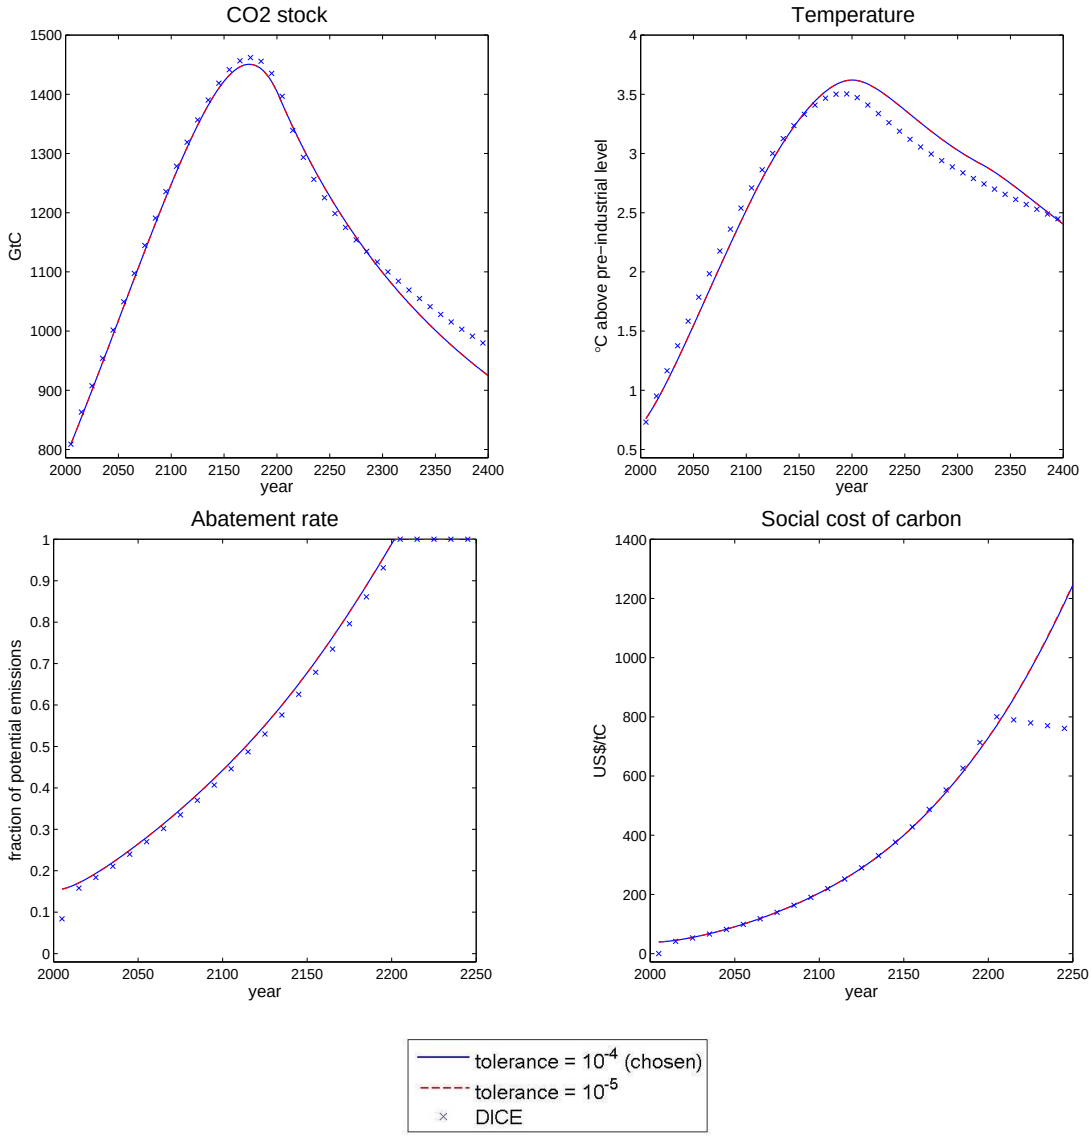

Figure 10 shows robustness of the results to a decrease in the convergence tolerance.

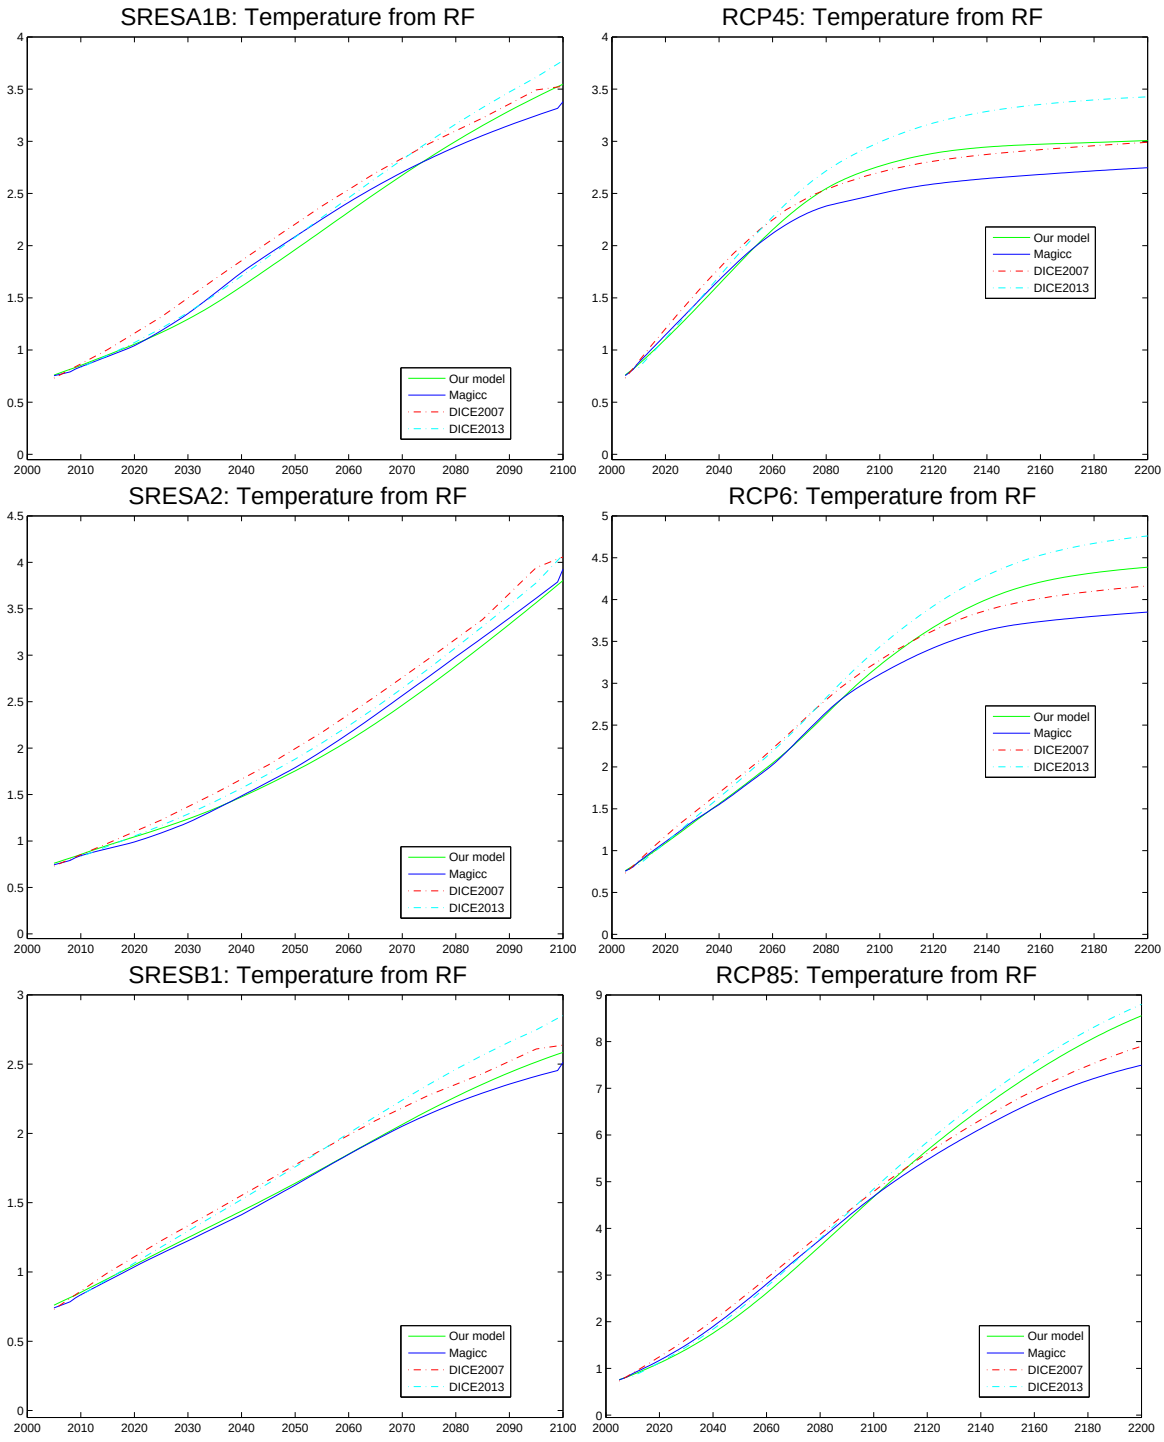

Figure 11 compares our simple calibration and those of the DICE 2007 and DICE 2013 to the MAGICC model. The scenarios are the same as in Figure 3, but here we directly compare the temperature response to the radiative forcing generated in the different scenarios by all greenhouse gases.

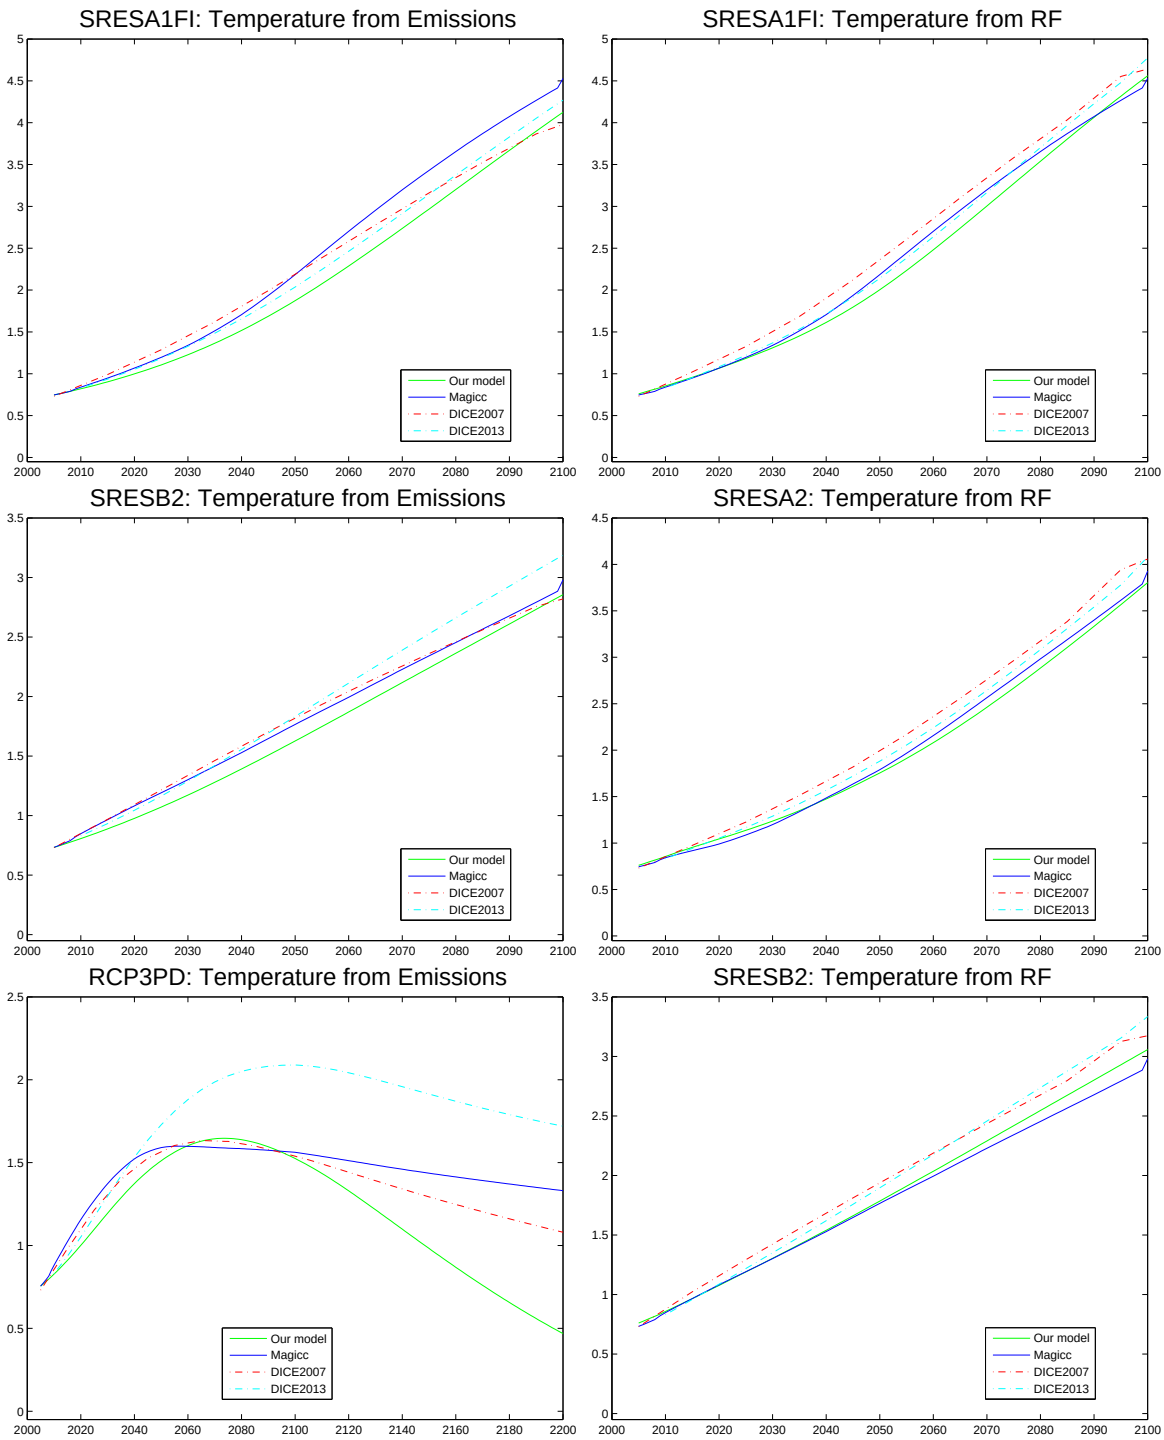

Figure 12 compares our simple calibration and those of the DICE 2007 and DICE 2013 to the MAGICC model. We show the temperature response to the remaining emission scenarios used in the 4^{th} (SRES) and 5^{th} (RCP) Assessment Reports of the IPCC. On the left we show the direct emission response, on the right the temperature response of the implied radiative forcing by all greenhouse gases.