

UC Berkeley

CUDARE Working Papers

Title

Risk and aversion in the integrated assessment of climate change

Permalink

<https://escholarship.org/uc/item/1562s275>

Authors

Crost, Benjamin
Traeger, Christian P.

Publication Date

2011-11-01

University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

Year 2011

Paper 1104R

Year 2010

Paper 1104

Risk and Aversion in the Integrated Assessment
of Climate Change

Benjamin Crost & Christian Traeger

Risk and Aversion in the Integrated Assessment of Climate Change

Benjamin Crost & Christian Traeger

Department of Agricultural & Resource Economics, UC Berkeley

CUDARE Working Paper 1104

Version of July 2011

!Print preferably in COLOR!

Abstract: We analyze the impact of damage uncertainty on optimal mitigation policies in the integrated assessment of climate change. Usually, these models analyze uncertainty by averaging deterministic paths. In contrast, we build a consistent model deriving optimal policy rules under persistent uncertainty. For this purpose, we construct a close relative of the DICE model in a recursive dynamic programming framework. Our recursive approach allows us to disentangle effects of risk, risk aversion, and aversion to intertemporal substitution. We analyze different ways how damage uncertainty can affect the DICE equations. We compare the optimal policies to those resulting from the wide-spread ex-ante uncertainty approach averaging deterministic paths.

JEL Codes: Q54, Q00, D90, C63

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

Corresponding Author:

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

There is a wide consensus that climate change takes place today and, even more so, that it will take place in the future. The precise changes and their socio-economic consequences remain uncertain at least for several more decades. This paper analyzes how uncertainty affects the social cost of carbon and the optimal abatement rate in the integrated assessment of climate change. Our main contribution is to disentangle the effects of risk, risk aversion, and intertemporal substitutability. Employing a replica of Nordhaus's (2008) DICE model, we also explain how seemingly innocuous details of the model imply that damage uncertainty can either cause a large increase or a small decrease of the optimal carbon tax.

We implement a slightly simplified version of Nordhaus's (2008) DICE model in a recursive dynamic programming framework. Our model features persistent uncertainty, an annual time step, and an infinite planning horizon. Optimal policies in a truly stochastic model depend on historic realization of the random variable. We discuss optimal expected trajectories and their robustness. We also present the optimal control rules as functions of the realized stock variables, capital, carbon, and temperature. Our discussion focuses on damage uncertainty. We show that the way in which uncertainty enters the damage function has a major effect on optimal mitigation policy.

We employ Epstein & Zin's (1989) recursive utility approach to distinguish a decision maker's propensity to smooth consumption over time from his Arrow-Pratt risk aversion. This step is important for several reasons. First, the approach enables us to disentangle whether observed differences between the certain and the uncertain setting are driven by risk, by risk aversion, or by the desire to smooth consumption over time. Second, evidence suggests that individual's tend to be more averse to risk than to intertemporal substitution. Third, also from a normative point of view there is no obvious argument that these a priori different dimensions of preference should coincide. Finally, we show that the disentangled model implies a significantly higher abatement rate and social cost of carbon than the standard model. For a preview of this result see Figure 1 on page 11.

Our approach to capturing uncertainty differs significantly from the wide-spread approach, often referred to as Monte-Carlo, which is a model of ex-ante uncertainty. Implicitly, this approach assumes that all uncertainty is resolved before the decision maker sets his policies. Nordhaus (2008) models the impact of ex-ante uncertainty on the optimal carbon tax. For this purpose, he draws selected parameters randomly from a set of distributions and runs the DICE model for each realization. He then interprets the average over the resulting deterministic paths as addressing uncertainty. Such a way of modeling uncertainty is an interesting first approximation. However, it is not a consistent approach to capturing uncertainty. In the ex-ante approach, the modeled decision maker has perfect knowledge about the true parameters by the time that he sets the policies and the model runs are initiated. In order to make statements about uncertainty, the approach averages over different model runs

in which different parameters govern the deterministic system. In contrast, under real uncertainty, the decision maker has to pick a joint policy for scenarios that differ only in future realizations of the random variable. The latter approach is computationally more demanding and, in an optimal control framework like used by Nordhaus (2008), quickly becomes infeasible as the decision tree branches out. In contrast, our recursive formulation permits us to introduce uncertainty in every period over an infinite planning horizon without major costs. We compare the resulting impact of uncertainty on the optimal policies to the ex-ante approach in order to analyze the quality of the ex-ante approximation. Ackerman, Stanton & Bueno (2010) recently complemented Nordhaus's (2008) uncertainty evaluation in DICE. Keeping the ex-ante uncertainty approach, they significantly increase the number of draws in their simulations. The authors assume uncertainty over climate sensitivity and the damage exponent. However, when comparing their results to Nordhaus (2008) they not only increase uncertainty, but also change parameter values. Their model does not translate the results of the simulations into changes of the optimal climate policy, like e.g. the optimal carbon tax.

Closest to our approach are the papers by Kelly & Kolstad (1999), Leach (2007), and Ha-Duong & Treich (2004). In their seminal paper Kelly & Kolstad (1999) implement a recursive version of DICE to analyze the effect of learning about climate sensitivity. An extended analysis is provided by Leach (2007). The main focus of these papers is learning time and how the possibility of learning changes the optimal policies. Compared to these learning oriented papers, we take a step back and focus on the analysis of uncertainty without parametric learning.¹ We take this step for three reasons. First, learning is a process that reduces uncertainty over time. To understand what drives the effect of learning, it is useful to have a thorough understanding of the effects of uncertainty. Even without analyzing how a decision maker reacts to different learning environments, we already compare over 30 different scenarios based on differences in preferences and in the way that damage uncertainty enters the equations. Second, in order to model parametric learning we would have to increase the state space of our model. Due to the curse of dimensionality, this increase would significantly limit our ability to do the type of comparisons undertaken in this paper, in quantity and in quality. Third, as long as temperature increases, we can justify the permanent nature of our uncertainty by the assumption that learning over damages at lower temperatures does not contain much information about damages at higher temperatures. For all modeled scenarios, temperatures are rising over the next one and a half centuries, which is the time-span most relevant for current policy. Moreover, when physical processes permit us to learn, Kelly & Kolstad (1999) and Leach (2007) find that learning time lies in the same order of magnitude. In contrast to Kelly & Kolstad (1999) and Leach (2007), we add the ability to disentangle the

¹In every period, the decision maker learns about the realizations of the damage parameters and the implied changes of the stock variables. However, we abstract away from a deeper learning where the decision maker learns about the distributions of the random variables, which would cause uncertainty to decrease over time.

effects of risk, risk aversion, and intertemporal substitutability. Moreover, we focus our analysis on various types of damage uncertainty, while the former papers model uncertainty about climate sensitivity.

Ha-Duong & Treich (2004) are the first to point out possible effects of disentangling risk aversion from intertemporal substitutability in relation to climate change. They build a simple numerical four period integrated assessment model incorporating Epstein & Zin (1989) preferences. The stochastic damage in their model is binary and investment is a fixed fraction of production.² The authors observe that increasing aversion to intertemporal substitution generally increases pollution, while increasing Arrow-Pratt risk aversion decreases pollution. They conclude that models that entangle these two a priori different preference characteristics tend to underestimate the effects of risk. In contrast, we use a full-fledged infinite horizon integrated assessment model and add a richer and more realistic uncertainty structure. We derive the actual magnitude of the effects of risk, risk aversion, and aversion to intertemporal substitution on the social cost of carbon and the optimal abatement rate. In addition to deriving optimal control rules, we compare results under the preference parameters used by Nordhaus (2008) to those based on preference estimates found in the asset pricing literature using the disentangled approach.

More remotely related to our analysis are the following papers. Keller, Bolker & Bradford (2004) model uncertainty about climate sensitivity and a climate threshold in an adapted version of DICE. They find that uncertainty decreases optimal abatement.³ Karp & Zhang (2006) discuss learning under uncertainty about climate sensitivity and marginal abatement cost in a stylized linear quadratic model. They find that learning generally decreases optimal abatement at a given information set. Richels, Manne & Wigley (2004) build ex-ante uncertainty about exogenous emissions, climate sensitivity, and temperature lag into their integrated assessment model called MERGE. They look at the effect of uncertainty under an exogenously given policy and do not optimize the control variables. Anthoff & Tol (2010) introduce ex-ante uncertainty about close to 1000 parameters in their integrated assessment model FUND and analyze the effect of uncertainty on exogenously given policies. Baker & Shittu (2007) give a survey of literature that incorporates uncertainty into the analysis of technical change in the climate change context. Weitzman (2009) argues that the multiplicative functional form in which damages usually enter integrated assessment models is somewhat arbitrary. He employs an abstract setting that captures damages and consumption in (sort of) a joint utility function. He finds that additive damages can significantly increase the willingness to pay for an emission reduction. We show that, under uncertainty, a much smaller transformation of the standard damage function can already imply significant differences in the evaluation.

Section 2 introduces the recursive utility specification and explains the disentan-

²In the model, an endogenous energy tax reduces energy input and then production, consumption, emissions, and damage.

³As the authors use a similar damage formulation as Nordhaus's (2008) our discussion also relates to their analysis.

gement of risk aversion from intertemporal substitutability. Section 3 introduces the climate enriched economy. Section 4 presents the results and section 5 concludes. Appendix A contains further graphs, while Appendix B completes the set of equations. The online Appendix C gives details on the numerical implementation and calibration of the model.

2 Welfare and Bellman equation

2.1 The Original Welfare Specification

This section introduces our welfare function and dynamic programming equation. The main feature that distinguishes our setting from other integrated assessment models is the ability to disentangle the effects of risk and risk attitude from effects driven by the desire to smooth consumption over time. For this purpose we employ a slightly modified version of the generalized isoelastic model introduced by Epstein & Zin (1989) and Weil (1990). In the standard intertemporally additive expected utility approach of the form $U = E \sum_t \exp[-\delta_u t] u(x_t)$ the curvature of the utility function u describes both, Arrow-Pratt risk aversion and intertemporal substitutability. However, evidence suggests that people tend to be more averse to substituting consumption into a risky state than into the certain future. Such behavior indicates that most people's Arrow Pratt risk aversion is higher than their aversion to intertemporal substitution (measured by the inverse of the intertemporal elasticity of substitution). In order to keep these two preference dimensions apart, we have to introduce two distinct aggregators, one for aggregating over time and another one for aggregating over risk. Time consistency then requires that uncertainty is evaluated at the point in time where it resolves, during a recursive aggregation of welfare over time. For more details see for example Traeger (2009). The corresponding Bellman equation for our setting becomes

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

The value function V represents the maximal welfare that can be obtained given the following state variables: Capital K_t , the stock of carbon in the atmosphere M_t , temperature T_t measured as deviation from preindustrial temperature, and time t . Utility within a period corresponds to the first term on the right hand side of the dynamic programming equation (1). It is a population L_t weighted power function of global per capita consumption C_t/L_t . The parameter η is the inverse of the intertemporal elasticity of substitution and captures the desire to smooth consumption over time (or aversion to intertemporal substitution). The second term on the right hand side

of equation (1) represents the maximally achievable welfare over the future. The term is discounted by the rate of pure time preference δ_u (also utility discount rate), which we assume to be 1.5%, in accordance with Nordhaus (2008). At any point in time t , the damages in the upcoming period and, thus, the capital stock K_{t+1} are uncertain. The parameter RRA characterizes the Arrow Pratt measure of relative risk aversion. Note that the standard model would implicitly assume $\eta = \text{RRA}$, in which case both exponents in the expression would disappear. The decision maker maximizes over the control variables C_t and the emission control rate μ_t .⁴ Equation (1) states that the value of an optimal consumption path starting in period t should be the maximal sum of the instantaneous utility gained in that period and the welfare gained from the continuation path. The control C_t balances the gratification of immediate consumption against the value of the future stocks. The control μ_t characterizes the abatement rate and balances consumption and capital investment against investment into the reduction of atmospheric carbon. We model the temperature T_t as an endogenous stock variable in order to capture the delay between an increase of the emission stock and an increase in global average temperature.

For a more detailed analysis of the interpretation of the parameters η and *RRA* see Epstein & Zin (1989) and Traeger (2010), who also derives the particular representation we employ here that is additive in the time step. Our baseline (N) is a scenario that employs Nordhaus's (2008) preference parameters $\eta = \text{RRA} = 2$. We compare this baseline to a scenario (D) based on disentangled preference estimates by Vissing-Jørgensen & Attanasio (2003), Basal & Yaron (2004), and Basal, Kiku & Yaron (2010) who build on Campbell's (1996) approach of log-linearizing the Euler equations in the asset pricing context. These papers suggest $\eta = \frac{2}{3}$ and find risk aversion coefficients around $\text{RRA} = 9.5$. For both scenarios we also run variations featuring risk neutrality $\text{RRA} = 0$ and extreme risk aversion $\text{RRA} = 50$.

2.2 A Conveniently Modified Bellman Equation

We find that the performance of our recursive numerical model improves significantly by expressing the relevant variables in effective labor terms and rewriting the Bellman equation accordingly. Exogenous technological progress is characterized by the variable

$$A_{t+1} = \exp[g_{A,t}]A_t \quad \text{with} \quad g_{A,t} = g_{A,t} \exp[-\delta_A t], \quad (2)$$

where $g_{A,0}$ denotes the initial growth rate, which thereafter declines exponentially over time. Similarly, population growth is captured as

$$L_{t+1} = \exp[g_{L,t}]L_t \quad \text{with} \quad g_{L,t} = \frac{g_L^*}{\frac{L_\infty}{L_\infty - L_0} \exp[g_L^* t] - 1}. \quad (3)$$

⁴In the numerical implementation of the model we generally maximize over the abatement cost Λ_t , which is a strictly monotonic transformation of μ_t (see 11). This switch of variables makes the constraints on the optimization problem linear.

Here L_0 denotes the initial and L_∞ the asymptotic population. The parameter g_L^* characterizes the convergence from initial to asymptotic population. Figure 9 in Appendix A shows the time paths of the exogenous variables, all of which are continuous time approximations to Nordhaus (2008). Expressing consumption and capital in effective labor terms results in the definitions $c_t = \frac{C_t}{A_t L_t}$ and $k_t = \frac{K_t}{A_t L_t}$. Using these definitions we transform equation (1) to

$$V(k_t A_t L_t, M_t, T_T, t) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} A_t^{1-\eta} L_t + \frac{\exp[-\delta_u]}{1-\eta} \exp[(1-\eta)g_{A,t}] A_t^{1-\eta} \times \exp[g_{L,t}] L_t \left(\mathbb{E} \left[(1-\eta) \frac{V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{A_{t+1}^{1-\eta} L_{t+1}} \right]^{\frac{1-\text{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\text{RRA}}}. \quad (4)$$

A further transformation maps the infinite time horizon conveniently onto the unit interval.⁵ For this purpose we introduce artificial time

$$\tau = 1 - \exp[-\zeta t] \in [0, 1] \quad (5)$$

and define

$$V^*(k_\tau, M_\tau, \tau) = \frac{V(K_t, M_t, t)}{A_t^{1-\eta} L_t} \Bigg|_{K_t=k_t A_t L_t, t=-\frac{\ln[1-\tau]}{\zeta}},$$

where A_t and L_t follow the exogenous time paths solving equations (2) and (3), with analytic solutions given in equations (17) and (18) in Appendix B. Under slight abuse of notation, we write variables x_τ to represent the variables $x_{\tau(t)}$ at time $\tau(t)$ satisfying equation (5). Dividing the fixed point equation (4) by $A_t^{1-\eta} L_t$ and rewriting in terms of V^* yields the new dynamic programming equation

$$V^*(k_\tau, M_\tau, T_\tau, \tau) = \max_{c_\tau, \mu_\tau} \frac{c_\tau^{1-\eta}}{1-\eta} + \frac{\exp[-\delta_u + g_{A,\tau}(1-\eta) + g_{L,\tau}]}{1-\eta} \times \left(\mathbb{E} [(1-\eta)V^*(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, T_{\tau+\Delta\tau}, \tau + \Delta\tau)]^{\frac{1-\text{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\text{RRA}}}. \quad (6)$$

The time step is now in artificial time.⁶ Our numerical simulations approximate V^* by Chebyshev polynomials and solve equation (6) for its fixed point by function iteration. Details are found in Appendix C. We recover the original value of a program from the value function V^* by the transformation $V(K_t, M_t, T_t, t) = V^*(K_t/A_t L_t, M_t, T_t, t) A_t^{1-\eta} L_t$. The marginal value of a ton of carbon is

⁵This time transformation also concentrates the Chebychev nodes on the near future, where most of the exogenously driven changes take place.

⁶Keeping a unit step in real time implies a time step $\Delta\tau = (1 - \exp[-\zeta])(1 - \tau)$ in artificial time and one period ahead artificial time becomes $\tau + \Delta\tau = 1 - [1 - \tau] \exp[-\zeta]$.

$\partial_{M_t} V(K_t, M_t, T_t, t) = \partial_{M_\tau} V^*(k_\tau, M_\tau, T_\tau, \tau) A_\tau^{1-\eta} L_\tau \Big|_{\tau=1-\exp[-\zeta t]}$ and the marginal value of a unit of the consumption-capital good is $\partial_{K_t} V(K_t, M_t, T_t, t) = \partial_{k_\tau} V^*(k_\tau, M_\tau, T_\tau, \tau) A_\tau^{1-\eta} L_\tau \partial_{K_\tau} K_\tau / A_\tau L_\tau \Big|_{\tau=1-\exp[-\zeta t]}$. The ratio of these two yields the social cost of carbon in current value units of the consumption-capital good:

$$SCC_t = \frac{\partial_{M_t} V}{\partial_{K_t} V} = \frac{\partial_{M_\tau} V^*}{\partial_{k_\tau} V^*} A_\tau L_\tau \Big|_{\tau=1-\exp[-\zeta t]} . \quad (7)$$

3 The Climate Economy

3.1 The model under certainty

The decision maker maximizes his value function under the constraints of the following stylized model of a climate enriched economy. The model is largely a reproduction of the DICE-2007 model by Nordhaus (2008). Except for our yearly (as opposed to decadal) time step, the two main differences are as follows. We do not model an explicit carbon cycle because state variables are computationally expensive in the recursive dynamic programming setting. Instead, we use an exogenously falling rate of carbon removal from the atmosphere. The exogenous decrease mimics that other carbon reservoirs, such as the ocean, reduce their uptake as they fill up over time. Online Appendix C provides details on how we calibrate this exogenous rate to DICE-2007. Similarly, instead of modeling ocean temperature explicitly, we introduce exogenous ocean cooling to mimick the endogenous cooling in DICE. Our calibrations yield a good match to the (optimal) DICE-2007 paths of CO₂ concentration and atmospheric temperature over the next four centuries and a very good match of the social cost of carbon and the optimal abatement rate for the next two centuries.⁷ In the following we focus on the most important equations and those that differ from DICE. Parameters are characterized and quantified in table A on page 38. The economy accumulates capital according to

$$k_{\tau+\Delta\tau} = [(1 - \delta_k) k_\tau + y_\tau - c_\tau] \exp[-(g_{A,\tau} + g_{L,\tau})] ,$$

where δ_K denotes the depreciation rate, $y_t = \frac{Y_t}{A_t L_t}$ denotes net production (net of abatement costs and climate damage) per effective labor, and c_t denotes aggregate global consumption of produced commodities per effective unit of labor. The removal rate of atmospheric CO₂ is $\delta_{M,t}$ falls exogenously over time to mimick the carbon cycle of DICE-2007 so that

$$\begin{aligned} \delta_{M,t} &= \delta_{M,\infty} + (\delta_{M,0} - \delta_{M,\infty}) \exp[-\delta_M^* t] \quad \text{and} \\ M_{\tau+\Delta\tau} &= M_{pre} + (M_\tau - M_{pre}) (1 - \delta_{M,\tau}) + E_\tau . \end{aligned} \quad (8)$$

⁷Note that our ocean cooling and decay rate are calibrated to the Nordhaus (N) preference scenario. Thus, we introduce an inertia that slightly reduces the relatively large effects of a change in preferences to the disentangled (D) scenario. The bias is negligible for the relatively small changes caused by uncertainty.

The pre-industrial emission stock M_{pre} is the steady state level in the absence of (annual) anthropogenic CO₂ emissions E_t . Anthropogenic emissions are the sum of industrial emissions and emissions from land use change and forestry B_τ

$$E_\tau = (1 - \mu_\tau) \sigma_\tau A_\tau L_\tau k_\tau^\kappa + B_\tau . \quad (9)$$

While emissions from land use change and forestry fall exponentially over time, industrial emissions are proportional to gross production $A_\tau L_\tau k_\tau^\kappa$ and reduced by abatement. Moreover, the DICE model includes an exogenously falling rate of decarbonization of production σ_t . For details see appendix B.

Net global GDP per effective unit of labor is obtained from the gross product per effective unit of labor as follows

$$y_\tau = \frac{1 - \Lambda(\mu_\tau)}{1 + D(T_\tau)} k_\tau^\kappa = \frac{1 - \Psi_\tau \mu_\tau^{a_2}}{(1 + b_1 T_\tau^{b_2})} k_\tau^\kappa \quad (10)$$

where

$$\Lambda(\mu_\tau) = \Psi_\tau \mu_\tau^{a_2} \quad (11)$$

characterizes abatement costs as percent of GDP depending on the emission control rate $\mu_t \in [0, 1]$. See appendix B for details on the coefficient Ψ_τ , which is assumed to fall over time. The percentage of world GDP that is lost to climate change related damages depends on the difference T_t between the current temperature and preindustrial temperatures and is characterized by

$$D(T_\tau) = b_1 T_\tau^{b_2} . \quad (12)$$

Nordhaus (2008) estimates $b_1 = 0.0028$ and $b_2 = 2$, which implies a quadratic damage function with a loss of 0.28% of global GDP at a 1°C warming. The temperature change is a delayed response to radiative forcing from CO₂ in the atmosphere

$$F_{\tau+\Delta\tau} = \eta_{forc} \frac{\ln \frac{M_{\tau+\Delta\tau}}{M_{preind}}}{\ln 2} + EF_\tau, \quad (13)$$

which is the sum of the forcing caused by atmospheric CO₂ and the non-CO₂ forcing that follows the exogenous process EF_τ described in appendix B. The temperature equation is

$$T_{\tau+\Delta\tau} = (1 - \sigma_{forc})T_\tau + \sigma_{forc} \frac{F_{\tau+\Delta\tau}}{\lambda} - \tilde{\sigma}_{ocean} \Delta T_\tau . \quad (14)$$

where $\sigma_{forc} = 3.2\%$ captures heat capacity related warming delay and $\sigma_{ocean} = 0.7\%$ relates to the cooling feedback caused by the temperature difference ΔT_τ between the atmosphere and the oceans (see Appendix B for a derivation). The values imply that a doubling of the CO₂ concentration relative to its preindustrial level will cause an equilibrium warming of 3°C (climate sensitivity).

3.2 Uncertainty

We analyze uncertainty over the parameters governing the damage function in equation (12). We follow Nordhaus (2008) in considering a normal distribution over the damage coefficient b_1 with a standard deviation of 0.13% (keeping the mean at 0.28%). For a mean warming of $3^\circ C$, the damage corresponding to the parameter's mean is 2.6% of world output, while a one sigma deviation around the mean corresponds to damages of 1.4% and 3.7% of world output. In addition we consider a normal distribution over the damage exponent b_2 . We chose a standard deviation of 0.35 (around the mean 2), which yields roughly the same damages as the b_1 uncertainty for plus/minus one standard deviation at a $3^\circ C$ temperature increase.⁸ A recent analysis by Hanemann (2009) contains a profound collection of arguments why actual damages from warming are underestimated in the DICE damage function. For our uncertainty analysis, we do not change the mean parameter values of the damage function, but we take studies like that of Hanemann (2009) as an indication that uncertainty is likely to be larger than expressed by the above standard deviations. Thus, we refer to the above specification as the “low” uncertainty scenarios and run our simulations as well for a set of “high” standard deviations. These high scenarios correspond to a standard deviation of 0.25% for the damage coefficient b_1 and a standard deviation of 0.5 for the damage exponent. These standard deviations translate into output losses of 4.8% (b_1) and a 4.4% (b_2) for a plus one sigma deviation and to losses of 0.3% (b_1) and 1.5% (b_2) for a minus one sigma deviation at a temperature increase of $3^\circ C$.

In scenarios with uncertainty over one damage parameter only, we calculate the value function optimizing over 5 Gauss-Legendre quadrature nodes per time step, approximating the normal distribution.⁹ Figure 15 in Appendix C shows that the results are not sensitive to using more quadrature nodes. In our scenarios with uncertainty on both parameters, we calculate the value function optimizing over 5×5 quadrature nodes per time step, approximating the bivariate normal distribution assuming independence between b_1 and b_2 uncertainty. Thus, the decision maker's control rules are based on a decision tree for which every subtree every year divides into 5 further subtrees (25 with joint uncertainty). The decision maker learns about the realizations of the random variables in every period and optimally adapts her behavior. In consequence, the most accurate way to represent the optimal policies is by means of control rules. However, time paths of the social cost of carbon and the abatement

⁸We calibrate slightly closer to the upper than to the lower one-sigma-deviation for a $3^\circ C$ increase. The damage for $b_2 = 2 - 0.35$ results in a 1.74% loss of GDP as opposed to 1.39% in the Nordhaus (2008) based b_1 uncertainty scenario and the damage for $2 + 0.35$ results in a 3.75% loss of GDP as opposed to 3.73% in the Nordhaus (2008) based b_1 uncertainty scenario. Note that for temperatures below $3^\circ C$ the impact of b_2 uncertainty decreases and vanishes for $T = 1^\circ C$ (as unity is unaffected by the exponent).

⁹Gauss-Legendre quadrature nodes make the discrete approximation to the integral over the distribution exact for the first $2n - 1 = 9$ moments. Neither our choice of a normal distribution nor our simulation model is meant to examine potential effects of fat tails.

rate are insightful and more suitable for the comparison of multiple scenarios. For the purpose of path comparisons, we will depict the particular time path obtained when the decision maker decides under uncertainty, but nature always happens to draw the expected value of the distribution. Later, we show that this path closely resembles the median path and the expected path obtained from sampling 10,000 runs with truly random draws by nature. We compare these time paths to those obtained from averaging deterministic paths, as done in the wide-spread ex-ante uncertainty approach discussed in the introduction. These simulations calculate the deterministic paths corresponding to each of our Gaussian quadrature nodes of the uncertain parameter and take the weighted average. Figure 16 in Appendix C shows that these simulations are not sensitive to the number of nodes either.

4 Results

Figure 1 presents the optimal abatement rate and the social cost of carbon (SCC) for the parameter choice of Nordhaus (2008) (N, dark blue) as well as for the assessment using disentangled preferences based on the estimates reviewed in section 2.1 (D, light green). The SCC corresponds to the optimal carbon tax.¹⁰ The solid lines represent assessment under certainty. The dashed lines introduce uncertainty over the damage coefficient b_1 , and the dotted lines represent uncertainty over the damage exponent b_2 . For the uncertain scenarios, the figure depicts the path where the decision maker chooses under uncertainty and nature happens to draw the expected value of the distribution in every period. Both, the abatement rate and the optimal carbon tax, are higher in the case of disentangled preferences. In numbers, the additional SCC over the next decade is between \$60 to \$90 higher in the disentangled D preference scenarios than in the N preference scenarios. Within the N-preference scenario, uncertainty and type of uncertainty make a difference that increases from \$4 to \$70 over the course of the century or, in relative terms, from 11% to 33%. In the disentangled scenario, the difference increases from \$11 to \$80 or 11% to 16% over the same time span. The abatement rate difference between N preferences and the disentangled approach increases from 10 percentage points in the present to about 25 percentage points at the end of the century. With D preferences, full abatement is optimal starting in year 130 if uncertainty over the damage exponent is considered and starting in year 140 otherwise (numbers rounded to the decade). With N preferences full abatement becomes optimal in year 180 when considering uncertainty of the damage exponent, and in year 200 otherwise. The full-abatement times can be observed in Figure 8 in Appendix A, where we depict the abatement rate as well as

¹⁰They differ in the case of full abatement. The usual optimal control framework calculates carbon taxes as the marginal value of the last unit abated from equation (11). In the case of full abatement, these correspond to the minimal necessary carbon tax that yields full abatement. In our dynamic programming approach, we obtain the value function of the problem, so that we can calculate the actual SCC also in the full abatement regime using equation (7).

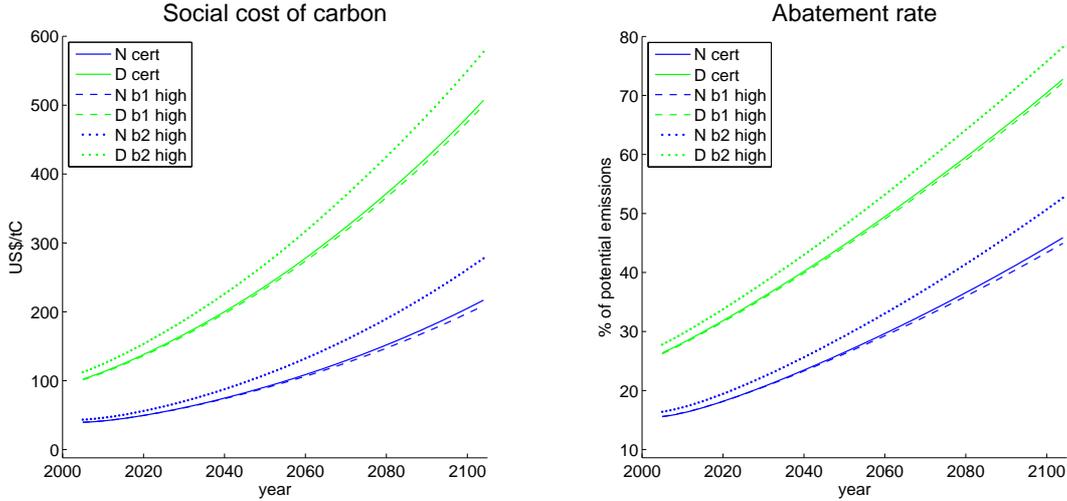


Figure 1 compares the optimal abatement rate and CO₂ tax based on the standard model with parameters $\eta = \text{RRA} = 2$ as in DICE-2007 (N, dark blue lines) with an assessment based on the estimates $\eta = \frac{2}{3}$ and $\text{RRA} = 9.5$ taken from the literature disentangling risk aversion from intertemporal substitutability (D, light green lines). “cert” denotes assessment under certainty (solid lines), “b1 high” introduces uncertainty over the damage coefficient (dashed lines), while “b3 high” introduces uncertainty over the damage exponent (dotted lines).

absolute emissions, temperature and damages for the different scenarios over the next two centuries. Insights from Figure 1 are, first, that uncertainty can have a significant impact on abatement and optimal taxes. Second, a distinction between risk attitude and the propensity to smooth consumption over time results in a major change of policy recommendations. And third, the type of uncertainty matters significantly and can even change the sign of the uncertainty effect. In the following we analyze these points in detail.

Figure 2 examines the dependence of the optimal abatement rate and the SCC on the degrees of risk aversion. The recursive preference setting allows us to vary risk aversion (RRA) while keeping aversion to intertemporal substitution (η) constant. The graph on the left fixes $\eta = 2$ as in DICE-2007, while the graph on the right fixes $\eta = \frac{2}{3}$ as suggested by the disentangling literature. In both scenarios uncertainty is on the damage exponent (high standard deviation). The graphs show the increase of the optimal SCC over a world of certainty. All graphs, including those corresponding to a risk neutral agent, imply a significant increase in the optimal SCC varying between \$5 and \$70 over the course of the century. However, the optimal policy hardly differs between the cases of risk neutrality, of $\text{RRA} = 2$ as in DICE-2007, and of $\text{RRA} = 9.5$ as suggested by the disentangled estimates. Only an extreme degree of risk aversion of $\text{RRA} = 50$ has an observable effect on policies, at least in the DICE case of $\eta = 2$. We conclude from Figure 2 that risk aversion itself does not induce significant effects on the optimal policies in our setting. Thus, the risk effects must be a consequence of other non-linearities in the underlying climate and economic system.

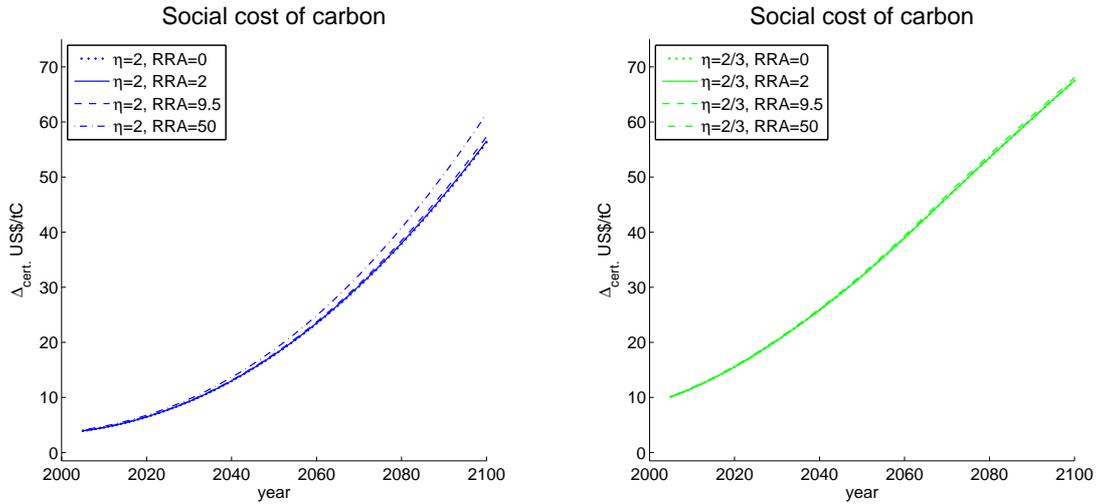


Figure 2 compares the optimal carbon tax for different degrees of risk aversion (RRA). The graphs show the increase Δ_{cert} of the SCC over the SCC under certainty. The left graph represents scenario N with $\eta = 2$ as in DICE-2007 (blue lines), while the right graph represents scenario D where $\eta = \frac{2}{3}$ as suggested in the disentangling literature (green lines). Uncertainty corresponds to a normal on the damage exponent (high standard deviation). In the case of scenario N extreme risk aversion is distinguishable from the other scenarios, while in the D scenario all lines lie on top of each other.

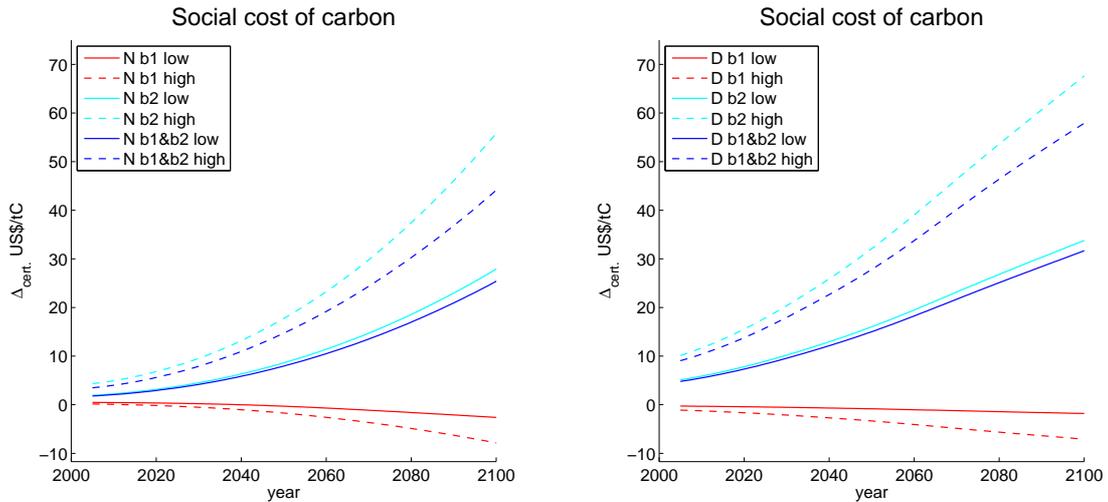


Figure 3 compares the optimal carbon tax for different types of uncertainty for scenario N (left) and for the disentangled scenario D (right). Uncertainty is over the damage coefficient “ b_1 ” and the damage exponent “ b_2 ”. The graphs depict the \$ increase over the optimal tax in the certain scenario.

Figure 3 compares the two different types of damage uncertainty for high and low variance and for the case of joint uncertainty. The graphs depict the increase of the SCC over the certain case for N preferences on the left and D preferences on the right. The most striking result depicted in Figure 3 is that uncertainty on the

damage coefficient b_1 reduces abatement and the SCC in both scenarios. Note that b_1 low (solid magenta) is precisely the type of damage uncertainty that Nordhaus (2008) adds to his ex-ante uncertainty version of the DICE model. While his normal distribution only reduces the SCC by a few dollars in the second half of the century, the dashed version featuring the higher standard deviation decreases the optimal carbon tax notably (and the optimal abatement rate by one percentage point at the end of the century). In contrast, the normal distribution over the damage exponent (cyan) significantly increases the SCC for both scenarios. The same is true for combined uncertainty on both parameters (dark blue lines). The increase in abatement corresponding to the depicted increase in the carbon tax by \$30 to \$60 at the end of the century corresponds to 3.5 to 7 percentage points. Recall that the standard deviations for the b_2 uncertainty scenarios were chosen to resemble the damages of the b_1 uncertainty scenario at a 3°C temperature increase and plus/minus one standard deviation shocks. Clearly it is not the magnitude of damage uncertainty that drives the difference between the two distinct uncertainty scenarios. The reason lies in the precise way in which uncertainty about the damage function enters the equations. The function that translates temperature related shocks into a loss of net production (and thus consumption and investment) is captured by equation (10), which can be rewritten as

$$Y_t^{net} = \frac{Y_t^*}{1 + b_1 T_t^{b_2}} =: f(b_1, b_2), \quad (15)$$

where $Y_t^* = (1 - \Lambda(\mu_\tau)) Y_T^{gross}$, i.e the total production less abatement expenditure. The uncertainty over b_1 corresponds to a linear variation in the denominator of equation (15). Thus, a convex function translates the variation in b_1 into variations of net GDP. Therefore, expected GDP under uncertainty over b_1 is actually higher than GDP using the expected coefficient.¹¹ The b_2 uncertainty also enters in the denominator of equation (15), but this time in the exponent of temperature. A straightforward calculation shows that the resulting transformation of b_2 into GDP is concave, so that expected GDP is lower than the GDP for the expected coefficient.¹² For the case in which both coefficients are uncertain, we find that the paths of abatement and the SCC lie between those of scenarios in which only one of the parameters is uncertain. Moreover, Figure 3 shows that the variance of the probability distribution that captures uncertainty is clearly relevant for determining optimal policies. Because uncertainty has a stronger effect when placed on b_2 , the size of this particular variance matters most. We also conclude that the way Nordhaus's (2008) DICE-2007 model places damage uncertainty onto b_1 gives rise to a somewhat special finding that

¹¹The second order derivative of the damage fraction in equation (15) with respect to b_1 is $\frac{\partial^2 f}{\partial b_1^2} = \frac{2T^{2b_2} Y^*}{(1+b_1 T^{b_2})^3} > 0$.

¹²The second order derivative of the damage fraction in equation (15) with respect to b_2 is $\frac{\partial^2 f}{\partial b_2^2} = -\frac{(1-2b_1 T^{b_2}) b_1 T^{b_2} (\ln[T])^2 Y^*}{(1+b_1 T^{b_2})^3} < 0$ as long as damages do not exceed half of the GDP.

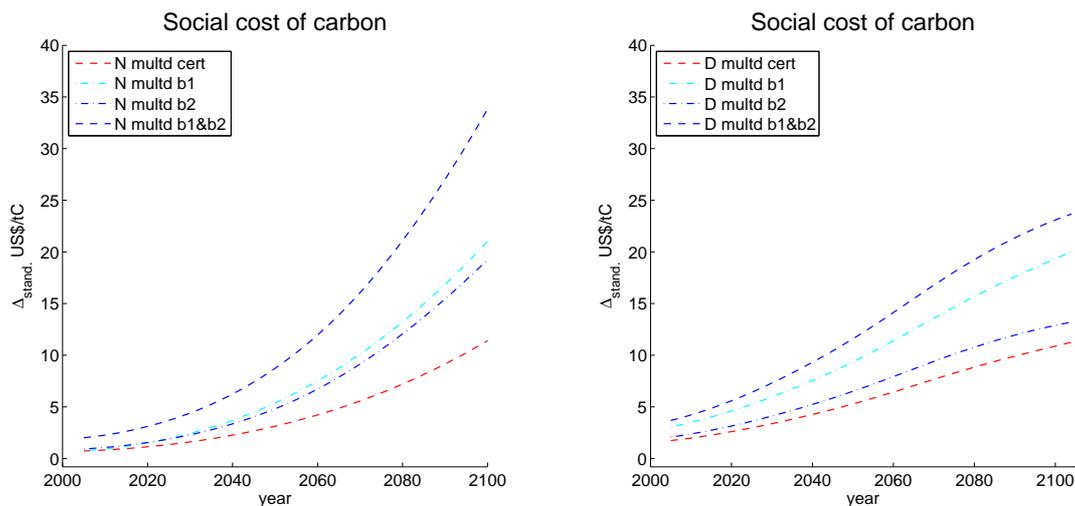


Figure 4 depicts the change in the optimal carbon tax when damages are linearly subtracted from gross output as in equation (16). The uncertainty scenarios are the high variance scenarios. The graphs present the \$ increase of the optimal tax within a given uncertainty scenario that results from changing equation (15) to equation (16). Also in the certain scenario the reformulation yields a higher optimal tax (magenta).

damage uncertainty is of little importance for optimal policies.

The non-linearities in the climate-economy interaction drive the observed effects of uncertainty on optimal policy. Therefore, the question arises why DICE places the GDP loss due to damages in the denominator, while the abatement expenditure enters in the numerator. In particular, the parameter b_2 is taken to be the fraction of GDP loss for a one degree Celsius increase over preindustrial temperatures. This interpretation relies on the approximation that $\frac{1}{1+\epsilon} = 1 - \epsilon$ for ϵ small. Figure 4 examines how good this approximation is for the magnitude of actual damages along the optimal path in DICE. All scenarios in Figure 4 correspond to the case of a high standard deviation. The figure depicts the difference between the SCC in the original damage formulation as in equation (15) and a formulation employing

$$Y_t^{\text{net}} = (1 - b_1 T_t^{b_2}) Y_t^* , \quad (16)$$

where b_1 directly characterizes the percentage loss of GDP caused by a one degree warming, not just an approximation. The SCC is notably higher for all scenarios. Even for the certain model runs, the SCC increases up to \$12 over the century, because the approximation in equation (15) implies a relatively smaller damage for high temperatures. Also note that the damage term is no longer convex in b_1 but simply linear. Thus, uncertainty over the coefficient b_1 alone has no effect at all. The depicted increase with respect to the earlier model formulation removes the convexity effect caused by the approximation in equation (15).

In our simulations the decision maker picks current consumption and abatement expenditure at the beginning of the period. Uncertainty about production then trans-

lates into the residual investment into next period's capital stock. This timing and uncertainty structure is implied by the usual formulation of the Bellman equation (1). However, such a setting allows the decision maker to use the capital stock in order to smooth consumption uncertainty, in particular, in our setting with iid shocks. While investment fluctuates significantly due to damage shocks, consumption might not be affected much by this volatility. However, some impacts of climate change like floods, droughts, hurricanes, or diseases also have a significant damage component that directly induces consumption volatility. To capture these effects we have to modify the Bellman equation. Let \bar{y}_τ denote expected production in the current period and $\Delta y_\tau = y_\tau - \bar{y}_\tau$ the random difference between actual and expected production. The decision maker now controls expected consumption q_τ . Actual consumption $c_\tau = q_\tau + \alpha \Delta y_\tau$ takes up the fraction $\alpha \in (0, 1)$ of production uncertainty. The remaining uncertainty translates into investment, changing the capital accumulation equation to the form

$$k_{\tau+\Delta\tau} = [(1 - \delta_k) k_\tau + \bar{y}_\tau - q_\tau + (1 - \alpha) \Delta y_\tau] \exp[-(g_{A,\tau} + g_{L,\tau})] .$$

For $\alpha = 1$ uncertainty fully translates into consumption. For $\alpha = 0$ we are back to the original setting. In order to solve this generalized problem, uncertainty aggregation in the dynamic programming equation now has to include current consumption. Therefore, we iterate over the expected value function, which we denote \bar{V} , and find the Bellman equation

$$\bar{V}(k_\tau, M_\tau, T_\tau, \tau) = \max_{q_\tau, \mu_\tau} \frac{1}{1 - \eta} \left[\mathbb{E} \left((q_\tau + \alpha \Delta y_\tau)^{1-\eta} + (1 - \eta) \times \bar{V}(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, T_{\tau+\Delta\tau}, \tau + \Delta\tau) \right)^{\frac{1-\text{RRA}}{1-\eta}} \right]^{\frac{1-\eta}{1-\text{RRA}}} .$$

The upper graphs in Figure 5 show the difference in the SCC with respect to the original formulation. The graph confirms that the two models coincide for $\alpha = 0$. A complete shift of uncertainty onto consumption increases the SCC slightly, by up to \$7 for N-preferences at the end of the century and less than \$1 in the disentangled approach. Splitting uncertainty between consumption and investment always yields a slightly higher SCC for the first part of the century as compared to pure investment uncertainty. However, in the second part of the century, a moderate shift of uncertainty into consumption can even decrease the SCC. In this case, optimal abatement is also reduced (lower left). This observation is explained by a combination of two effects. First, under consumption uncertainty, the decision maker is more responsive to risk because he is no longer in a position to fully smooth the consumption shocks. By reducing emissions he can reduce future risk. However, second, the decision maker also has an incentive to reduce his spending on abatement in order to increase his baseline consumption, cushioning the consumption shocks. In the near future, temperature increase and immediate damages are low and the decision maker abates more

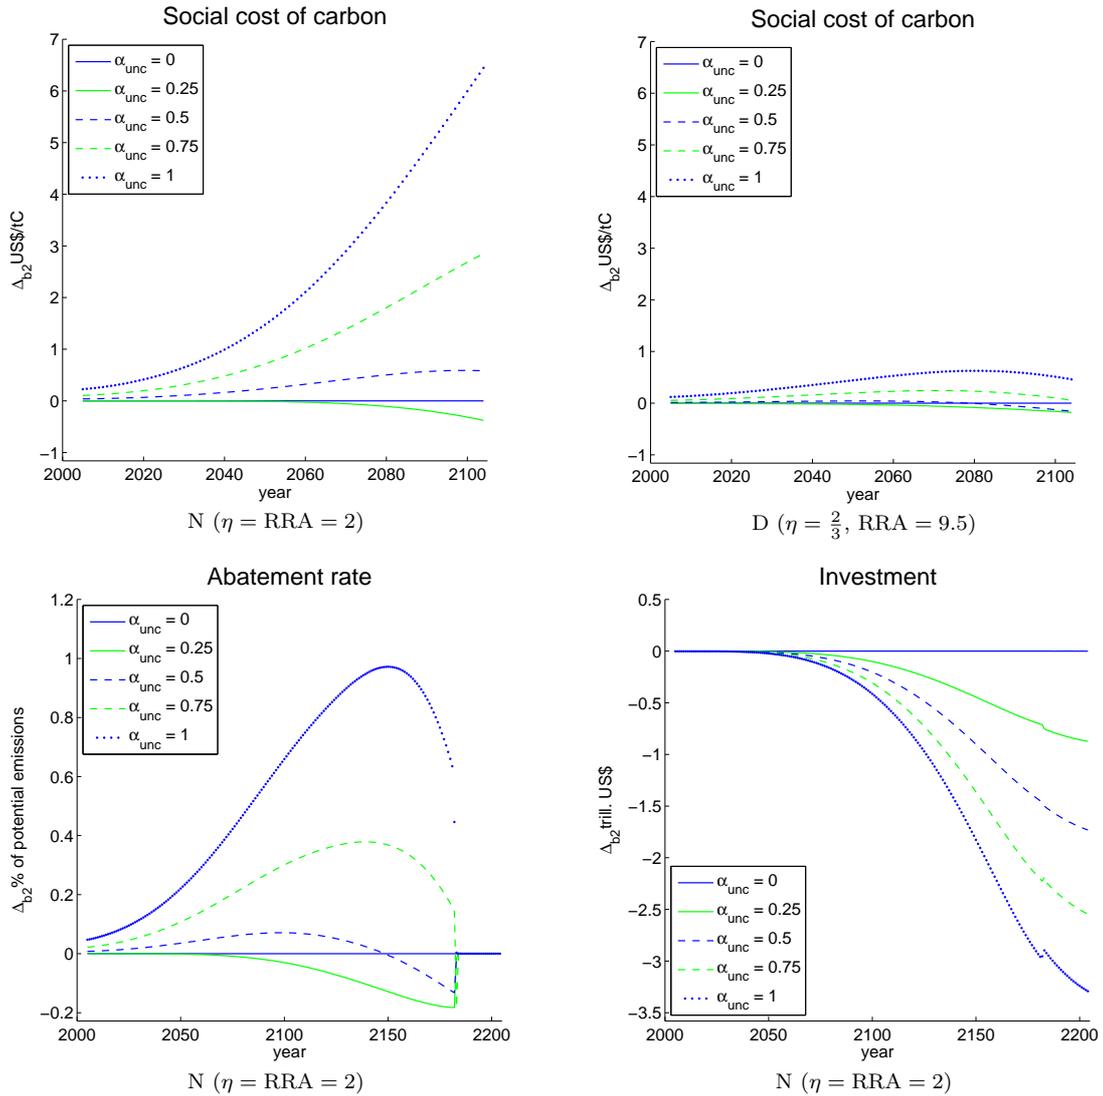


Figure 5 analyzes how the controls change in response to shifting uncertainty from investment ($\alpha = 0$) over to consumption ($\alpha = 1$). Uncertainty affects the damage exponent b_2 with high variance. The graphs present the increase over the original formulation where only residual investment is volatile. The parameter α_{unc} characterizes the fraction of consumption that covaries with the uncertain net production Y^{net} . The remaining uncertainty is swallowed by investment.

in order to reduce future risk. When temperature related shocks become large, an increasing amount of output is used to increase expected consumption. The graph in the lower left of Figure 5 shows that, far enough into the future for high enough damages, all scenarios exhibit a decrease in the additional abatement rate implied by consumption volatility (and for some of them the additional abatement is negative). Finally, the graph in the lower right of Figure 5 shows that investment continuously decreases in the shift from investment to consumption uncertainty. While this finding

might look surprising at first sight, it simply confirms that more expenditure goes into either abatement (long-term risk reduction) or consumption (immediate uncertainty cushion). The graphs depicted in Figure 5 change only marginally for the case of higher coefficients of relative risk aversion. Overall, we conclude that the comparison between a consumption and investment impact of uncertainty shows some interesting differences in the optimal policy, but the magnitude of the difference is small.

So far, our analysis compared optimal policy paths that were generated under the assumption that the decision maker optimizes under uncertainty, while nature happens to draw expected outcomes in all periods. In the following we discuss how our chosen path representation (expected draws) relates to other time path representations under uncertainty. For this purpose, we randomly draw the parameter realizations in every period for 10,000 time paths in the N and the D scenarios with damage exponent uncertainty (high variance). Note that the decision maker reacts optimally to each draw. Figure 6 depicts the resulting mean and median paths as well as the boundaries of the interval containing 95% of the individual paths at every point in time. In all of the graphs, the mean and the median of the simulations are virtually indistinguishable from the path generated by expected draws. The simulated 95% confidence interval is small for all observed quantities, except for damages themselves. Here, variations are sizeable and e.g. in the year 2100 the boundaries correspond to a damage of 0.7% of GDP for low years and of 4.1% of GDP for high years. The large effect of damage realizations in individual years on net world output also reflects in the investment paths, but is smoothed out in aggregate consumption (not shown). The grey lines depict 100 randomly selected paths.¹³

Figure 6 also compares our stochastic dynamic programming results to the ex-ante uncertainty approach discussed in the introduction. This wide-spread uncertainty simulation, or uncertainty proxy, simply averages deterministic paths. For our ex-ante simulation we use the same Gauss-Legendre quadrature nodes on the damage exponent as before. However, we ‘draw’ the parameter only once and the decision maker optimizes a deterministic program given the particular damage function realization. After doing so for each quadrature node, we take the weighted average resulting in the blue line. Figure 6 in Appendix C shows that the path is not sensitive to the number of quadrature nodes we use. For N-preferences, the ex-ante approach gives a good approximation to the stochastic simulations for the economic and climate variables in the near future. During the second half of the century, when temperatures have risen more significantly, the ex-ante approach underestimates the optimal cost of carbon and abatement rate more notably, by about \$30 and 6 percentage points at the end of the century. In contrast, the ex-ante approach yields an almost perfect match in terms of the expected damage simulation. The most crucial divergence of the ex-ante simulations happens for the carbon concentrations. For

¹³The highest spikes in the damage graph are truncated in order to have a better resolution. The highest spike in the 100 path sub-sample happens in the second half of the next century and takes away 37% of the production.

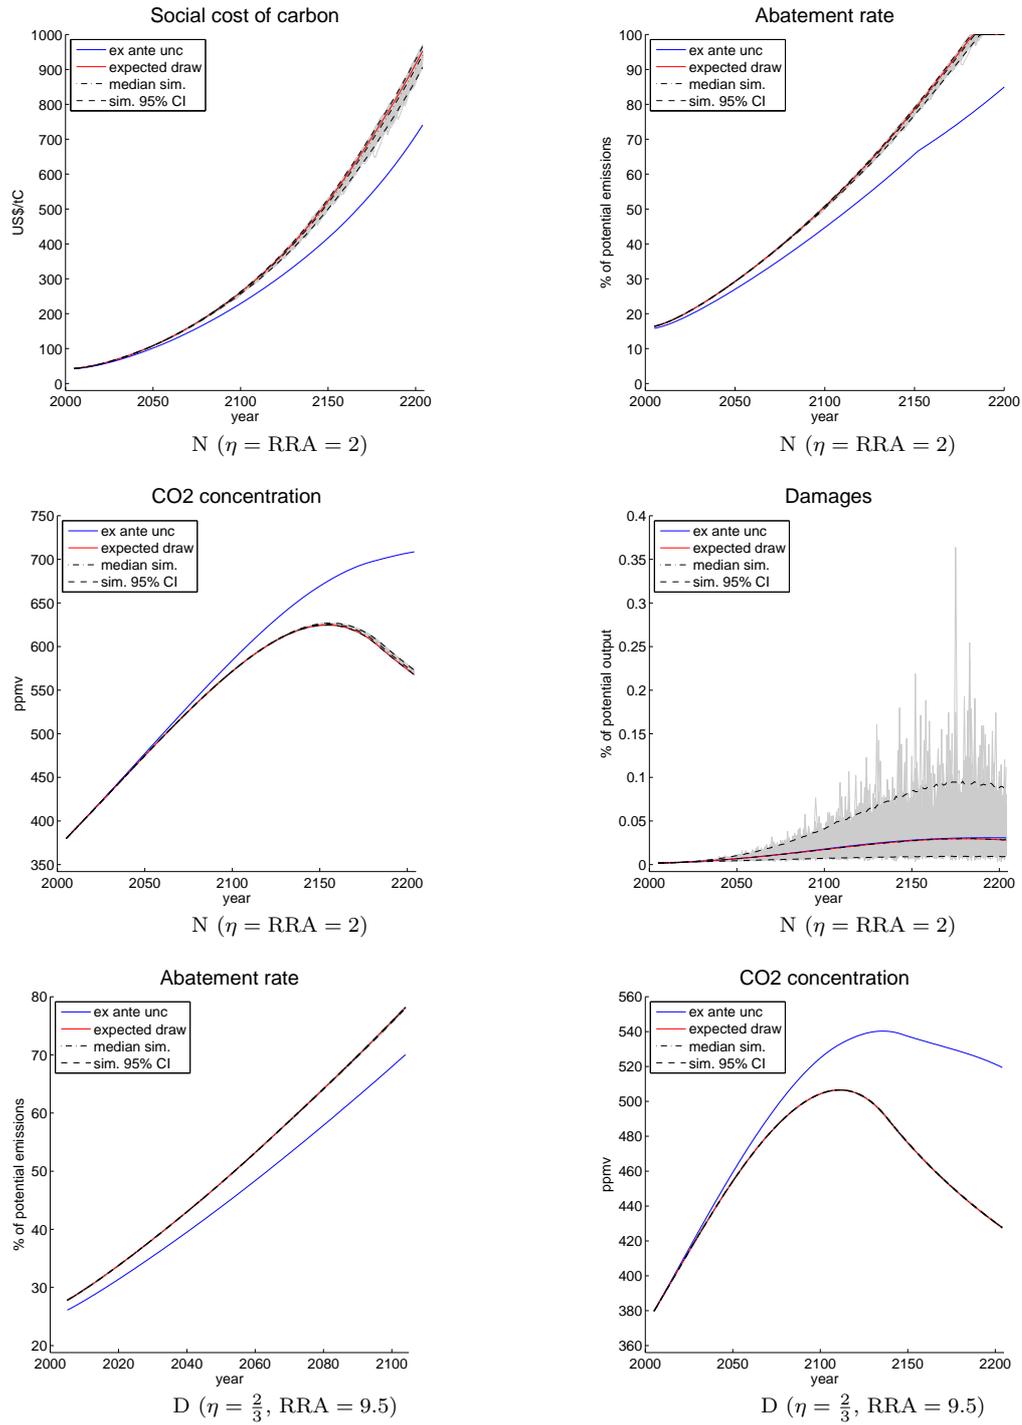


Figure 6 compares different representations of the expected time development under uncertainty over the damage exponent (high variance). The figure also compares our stochastic programming results to an ex-ante uncertainty approach averaging deterministic paths (blue). The grey lines represent 100 randomly selected paths out of the 10,000 random paths underlying the mean, median, and simulated 95% confidence interval.

N-preferences, the optimal peak level of carbon is about 100 ppm higher than in the stochastic programming simulation. The two lower graphs in Figure 6 present the optimal abatement rate and the CO₂ concentrations for disentangled preferences. Here, already the current policy recommendation that derives from the ex-ante approach underestimates the optimal abatement rate by over 2 percentage points (almost 10%). The optimal peak level of carbon would be overestimated by about 50 ppm. In both ex-ante simulations, N and D, the reduction of atmospheric CO₂ concentrations is also pushed further into the future. Overall, we conclude that the ex-ante simulation yields a decent approximation to optimal control rules in the near future. However, in the longer term, the ex-ante results differ more notably from the stochastic model with persistent uncertainty. In particular, when the policy parameter of interest is the optimal peak CO₂ level, the results of ex-ante simulations should be taken with care.

Any individual path under uncertainty is of Lebesgue measure zero. Therefore, the most accurate representation of optimal policies is by representing the control rules. The control rules characterize the optimal policy conditional on the current levels of the stock variables, which are determined by historic shocks as well as historic policy. Figure 7 depicts some of the optimal control rules for the year 2020. The blue dots mark the 2020 value along the expected path. The two upper graphs present the optimal abatement rate under certainty for the N preference structure used by Nordhaus (2008). The graph to the left presents the abatement rate in response to changes in the CO₂ concentration and in temperature. Capital is held fixed at the expected 2020 level. We observe that a higher global temperature leads to an almost linear increase in the optimal abatement rate. An increase in the concentration of CO₂ from pre-industrial levels to the level predicted under the optimal policy leads to a notable increase in the the optimal abatement rate. However, between 400-500ppm the increase of abatement in response to higher CO₂ concentrations starts flattening out. This flattening is a response to the logarithmic relation between CO₂ concentration and radiative forcing (see equation 13). The graph in the upper right depicts the optimal abatement rate in response to changes in the capital stock and in time, which is a proxy for the exogenous processes in DICE. The graph fixes temperature and CO₂ levels at the expected 2020 values. The abatement rate increases in both, capital stock and time. Keeping in mind that temperature and CO₂ levels are fixed, the interpretation of the abatement increase over time is mostly a response to the exogenously decreasing abatement cost (see Figure 9 in Appendix A for an overview of the exogenous processes in DICE). The graph in the middle left of Figure 7 shows the SCC, or optimal carbon tax, corresponding to the abatement rate shown in the first graph on the upper left. It displays the same pattern, including the decreasing responsiveness to increases in the CO₂ concentration. For very high temperatures, the marginal value of avoided emissions even decreases in the carbon concentration, because the temperature response to emission changes decreases for high carbon stocks. Currently, global policy is not following the optimal path and, most likely, it will

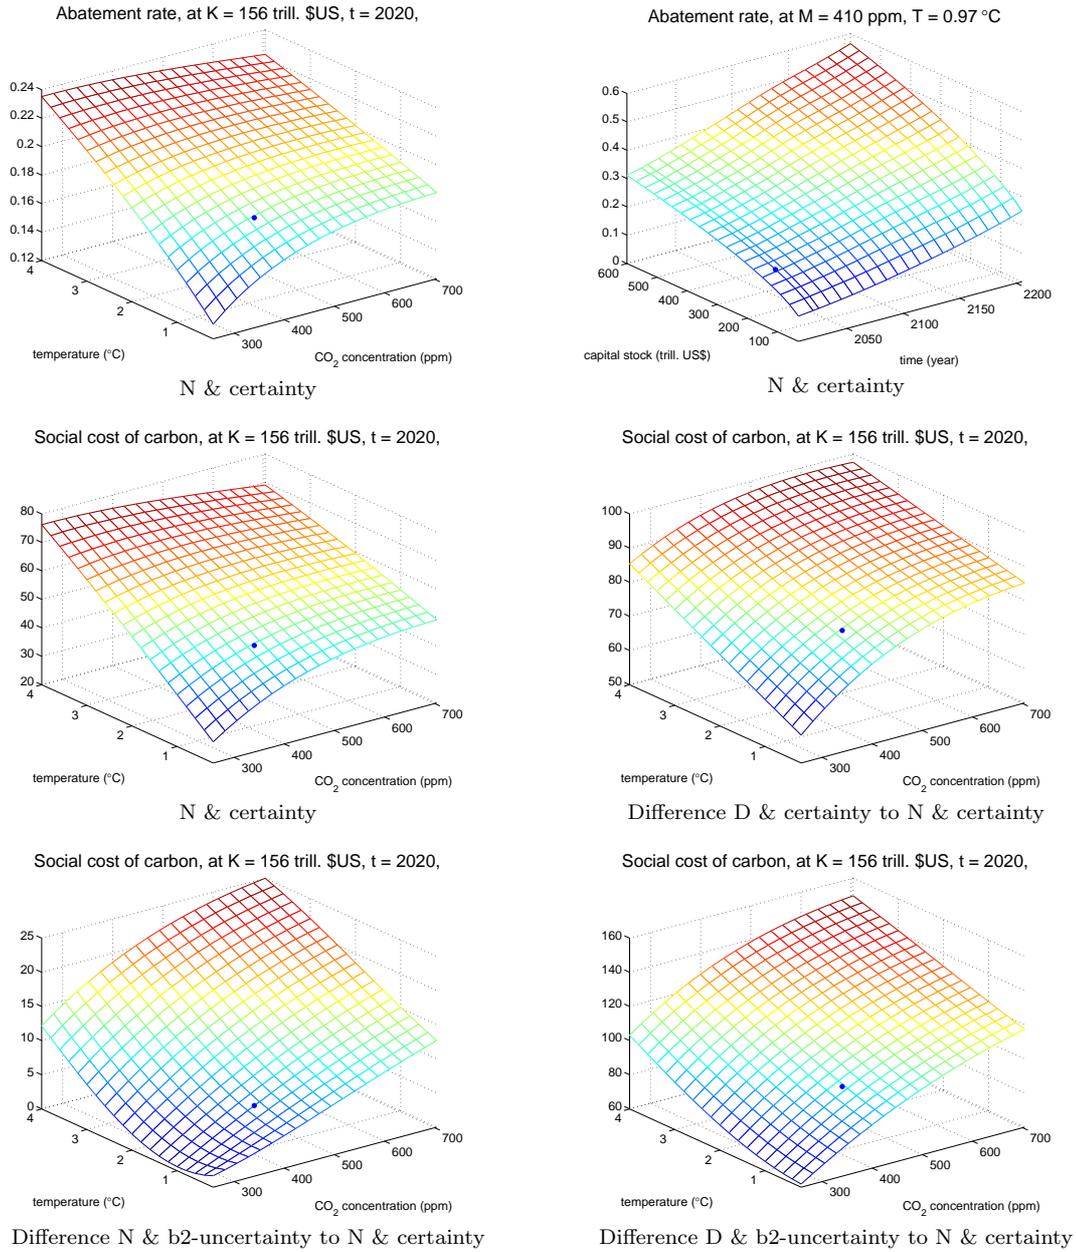


Figure 7 presents a selection of optimal control rules for the year 2020. The first three graphs show optimal abatement and the SCC based on a deterministic run employing the original DICE-2007 preferences (N). The graph in the middle right shows the optimal increment of the carbon tax for disentangled preferences. The graphs in the lower line show the increase in the social cost of carbon over the deterministic N model caused by introducing uncertainty over the damage exponent. The left graph depicts this increase for uncertainty and entangled preferences (N), and the right graph for uncertainty and disentangled preferences (D). The markers (blue dots) show the optimal policies along the deterministic path with DICE-2007 preferences. These values are also used to fix the states that are not shown.

not do so in the near future either. Hence, the actually optimal carbon tax in the year 2020 in a deterministic world will lie at a higher point corresponding to a higher carbon level to the right of the blue marker.

The graph in the middle right of Figure 7 depicts the differences in the optimal carbon tax between the D and the N scenarios under certainty. The interpretation of this difference is the optimal increase in the carbon tax when taking account of disentangling the propensity to smooth consumption from risk aversion. The blue dot at about \$80 marks the expected values when policy historically followed the optimal deterministic path with N-preferences. In case we follow the D-optimal path, carbon concentrations and the difference in the optimal carbon tax would be slightly lower. Most likely, the actual carbon stock in 2020 will be higher, and the optimal increment of the SCC will exceed \$80. The graphs in the lower line of Figure 7 show the change in the SCC under uncertainty. The graph on the left shows the effect of uncertainty over the damage coefficient under N preferences. The graph on the right evaluates the same uncertainty for D preferences. Both graphs show the optimal increment of the carbon tax over the values under certainty and N-preferences (i.e. over the values in the graph on the middle left). The dent surrounding a 1°C temperature increase, better observable in the left graph, is due to damage coefficient insensitivity at a unit temperature increase. The markers once more refer to the expected values under N preferences and certainty. The carbon tax should increase about \$ 8 if the policy maker decides to take account of damage coefficient uncertainty in year 2020. It should increase more if the preceding policy has not followed the optimal path. The optimal tax should increase about \$ 95 over the deterministic N-preference case, if the policy maker additionally realizes that he should disentangle risk attitude from the willingness to substitute consumption over time. Given current policy will yield higher than optimal carbon stocks in 2020, the optimal increment will most likely exceed \$ 100. These increments approximately triple the SCC in the year 2020 as compared to the deterministic standard model. Finally, note that the uncertainty increment does not flatten out as much in the CO₂ concentrations as the deterministic control rules.

5 Conclusions

The paper incorporates damage uncertainty into a truly stochastic integrated assessment model. Our recursive model allows us to disentangle the effects of risk, risk aversion, and aversion to intertemporal substitution. We find that risk has a notable impact on optimal policies. However, it is not the decision maker's risk aversion that causes the policy impact of risk. It is a consequence of the nonlinearity of the equations translating climate change into production losses. We have quantified how different functional formulations of damage uncertainty affect the social cost of carbon and the optimal abatement rate. Uncertainty about the damage coefficient results in a small decrease of the optimal abatement rate, an effect that we traced to back

to an approximation used in DICE. In contrast, uncertainty about the damage exponent results in a sizeable increase in abatement. We also quantified how the optimal policy differs between situations where uncertainty directly affects consumption and situations in which capital can be used to smooth consumption volatility. In general, these uncertainty effects imply a moderate increase of the optimal mitigation level and the social cost of carbon in the near future. As CO₂ concentrations and temperatures rise, the effect of uncertainty on optimal policy increases. Optimal policies are sensitive to the precise way uncertainty is incorporated into the model. This sensitivity emphasizes the importance of estimating precise functional forms of the damage function, including a careful functional identification of the remaining uncertainty.

A larger correction to optimal policy stems from disentangling risk aversion from the propensity to smooth over time. Recursive preference estimates find that risk aversion is higher and the propensity to smooth consumption over time is lower than the parameter describing both characteristics jointly in the standard model and in DICE. While the increase in risk aversion only has a very minor effect, the decrease in aversion to intertemporal substitution has a major effect on optimal policies. This latter parameter change more than doubles the current social cost of carbon and significantly increases optimal abatement, by increasing the willingness to invest into the future. Within the standard model, parameter changes that increase the willingness to invest into the future have generally been rejected from a positive perspective, arguing that such a model would not match observed interest rates. The current parameter change, in contrast, is a result of using a model that is successfully employed to better explain observed interest rates, including the risk free rate puzzle and the equity premium puzzle.

We compare our result from a truly stochastic integrated assessment model to a wide-spread method of averaging deterministic paths (ex-ante approach). We find that current policies are predicted quite well by the simplified ex-ante approach, in particular for the preference specification used in the original DICE model. The ex-ante approach is less convincing when it comes to future optimal policies and state variables. In general, the ex-ante approach underestimates the social cost of carbon and optimal abatement rate in our simulations. The approximation is worse in a model that corrects the intertemporal elasticity of substitution for the findings from the disentangling literature. The ex-ante approach was most notably off for optimal carbon peak concentrations. Here, an ex-ante simulation with standard preference parameters and uncertainty over the damage coefficient overestimates the optimal CO₂ peak level by about 100ppm. Again, the error of the ex-ante method will depend on the particular functional form in which uncertainty enters the equation. A clear message is that ex-ante Monte Carlo simulations should be taken with care when discussing long term objectives like the peak levels of CO₂ or temperature, which are extremely prominent in the policy debate.

A further advantage of our dynamic programming structure is that we derive the optimal control rules rather than just expected policy or temperature paths. These

control rules show how a policy maker reacts optimally to given states, which most importantly include the CO₂ concentration, but also temperature and capital stock. For example, we show how the optimal policy in the year 2020 varies depending on the actual carbon level at the beginning of year 2020. Solving large integrated assessment models by dynamic programming is difficult and the cost of additional state variables is high. Therefore, our solution method will not replace deterministic or ex-ante reasoning in general. However, our approach shows that risk aversion has a negligible impact on optimal policy, while aversion to intertemporal substitution has a major impact. The standard model cannot disentangle the two parameters and has to choose a single value that captures both. Our findings suggest that integrated assessment modelers using the standard model should set the parameter to match aversion to intertemporal substitution. Aversion to intertemporal substitution is lower than the value of the parameter jointly used in most integrated assessment models. Adjusting the parameter increases the weight given to the long-run consequences of climate change, increasing abatement. Opposed to the intuition derived from the standard model, this parameter change should not be accompanied by an increase in pure time preference. In explaining observed interest rates, including the equity premium puzzle, the disentangled model finds a rate of pure time preference that is generally even lower than the 1.5% that we adopted from DICE-2007.

Appendix

A Additional Graphs

Figure 8 shows a comparison of different scenarios for a 200 year time horizon. The scenarios are the same as in Figure 1, distinguishing the original DICE-2007 preference structure (blue) from the disentangled approach (green). Once more solid lines represent certainty, dashed lines high variance on the damage coefficient, and dotted lines high variance on the damage exponent. In addition to the policy parameters we also compare expected damages, CO₂ concentration, and temperature. Figure 9 shows the exogenous processes underlying the DICE-2007 model.

B Additional Equations Characterizing Climate and Economy

Where possible, we replace the difference equations of the DICE-2007 model by their solutions in continuous time. Table A presents the parameter values. The exogenous equations of motions for labor productivity A_t and population L_t summarized in

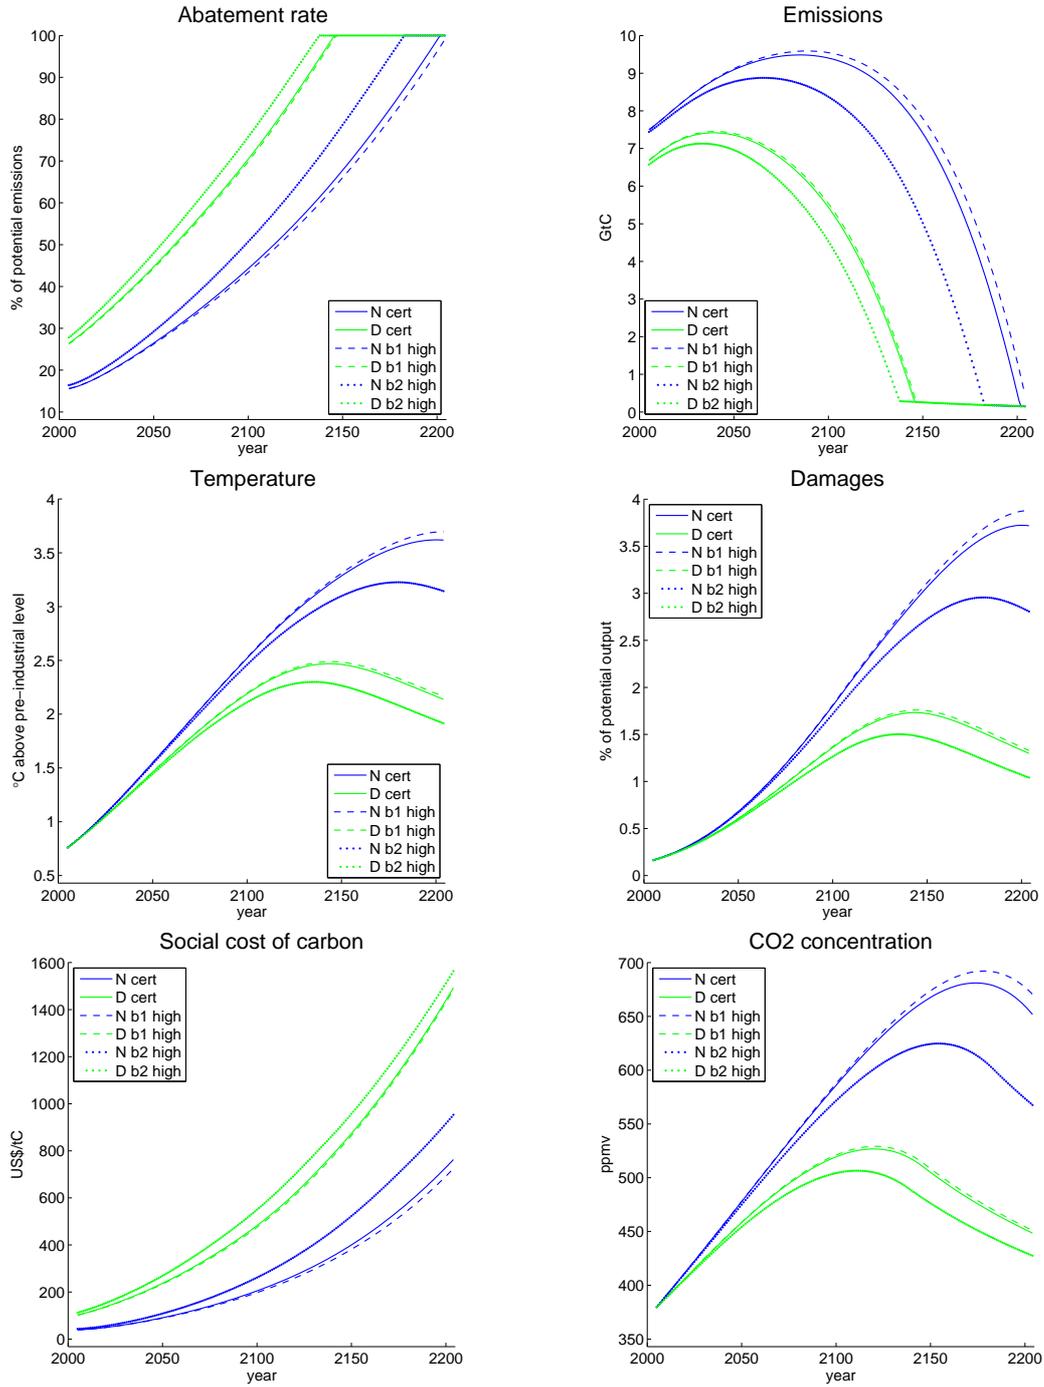


Figure 8 shows the time paths of the abatement rate, emissions, temperature, and damages for the N scenario with parameters $\eta = RRA = 2$ as in DICE-2007 (dark blue lines) and the disentangled D scenario with $\eta = \frac{2}{3}$ and $RRA = 9.5$. “cert” denotes assessment under certainty (solid lines), “b1 high” introduces uncertainty over the damage coefficient(dashed lines), while “b3 high” introduces uncertainty over the damage exponent (dotted lines).

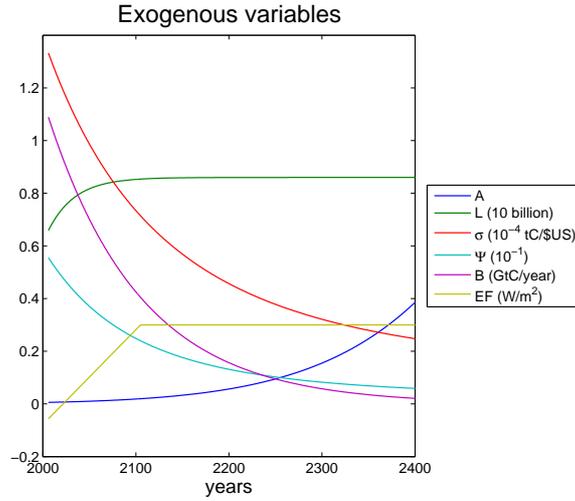


Figure 9 shows the time paths of the model's exogenous parameters: labor productivity (A), population (L), rate of decarbonization (σ), abatement cost coefficient (Ψ), emissions from land-use change and forestry (B), and exogenous forcing (EF).

equations (2) and (3) have the analytic solutions

$$A_t = A_0 \exp \left[g_{A,0} \frac{1 - \exp[-\delta_A t]}{\delta_A} \right] \quad \text{and} \quad (17)$$

$$L_t = L_0 + (L_\infty - L_0)(1 - \exp[-\delta_L^* t]) . \quad (18)$$

Emissions in equation (9) are subject to an exogenous baseline decarbonization of production

$$\sigma_t = \sigma_{t-1} \exp[g_{\sigma,t}] \quad \text{with} \quad g_{\sigma,t} = g_{\sigma,0} \exp[-\delta_\sigma t] ,$$

adopting the DICE assumption of an exponentially declining rate of decarbonization $g_{\sigma,t}$. we also follow DICE in assuming an exponential decline of emissions from land use change an forestry,

$$B_t = B_0 \exp[g_B t] .$$

Making use of these equations we can rewrite the emissions equation (9) as

$$E_\tau = (1 - \mu_\tau) \left[L_0 + (L_\infty - L_0) \left(1 - [1 - \tau]^{-\frac{g_L^*}{\zeta}} \right) \right] \sigma_0 A_0 k_\tau^\kappa \exp \left[g_A \frac{1 - (1 - \tau)^{\frac{\delta_A}{\zeta}}}{\delta_A} + g_\sigma \frac{1 - (1 - \tau)^{\frac{\delta_\sigma}{\zeta}}}{\delta_\sigma} \right] + B_0 [1 - \tau]^{-\frac{g_B}{\zeta}} .$$

The coefficient of the abatement cost function Ψ_τ in equation (11) is assumed to follow

$$\Psi_t = \frac{\sigma_t}{a_2} a_0 \left(1 - \frac{(1 - \exp[g_\Psi t])}{a_1} \right), \quad (19)$$

with a_0 denoting the initial cost of the backstop (i.e. in 2005), a_1 denoting the ratio of initial over final backstop,¹⁴ and a_2 denoting the cost exponent. The rate g_Ψ describes the convergence from the initial to the final cost of the backstop. In artificial time equation (19) translates into

$$\Psi_\tau = \frac{\sigma_0 a_0}{a_2} \exp \left[g_\sigma \frac{1 - (1 - \tau)^{\frac{\delta_\gamma}{\zeta}}}{\delta_\gamma} \right] \left(1 - a_1 \left(1 - [1 - \tau]^{-\frac{g_\Psi}{\zeta}} \right) \right). \quad (20)$$

Exogenous forcing in equation (13) from non-CO2 greenhouse gases, aerosols and other processes is assumed to follow the process

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

Note that it starts out slightly negatively.

In the following we explain our approximations that mimic the temperature delay equations of the original DICE model. In DICE-2007 a system of equations that can be transformed to

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

governs atmospheric and oceanic temperatures. The parameter $\lambda = \frac{\eta_{forc}}{s}$ relates equilibrium forcing of a CO₂ doubling to the corresponding equilibrium temperature increase s (climate sensitivity). We can transform the equation into

$$\begin{aligned} T_{t+10} &= (1 - \sigma_{forc}^{dec} - \sigma_{ocean}^{dec})T_t + \sigma_{forc}^{dec} \frac{F_{t+1}}{\lambda} + \sigma_{ocean}^{dec} T_t^{Ocean} \\ &= (1 - \sigma_{forc}^{dec} - \sigma_{ocean}^{dec})T_t + \sigma_{forc}^{dec} \frac{F_{t+1}}{\lambda} + \sigma_{ocean}^{dec} [(1 - C4)(T_{t-10}^{Ocean} - T_{t-10}) + T_{t-10}], \end{aligned}$$

where $\sigma_{forc}^{dec} = C1 * \lambda \approx 0.28$ and $\sigma_{ocean}^{dec} = C1 * C3 \approx 0.073$ are decadal lag parameters governing how atmospheric temperatures adjust to radiative forcing and to oceanic temperature. We approximate temperatures in $t-10$ in the square bracket by temperatures in t , neglecting terms of second order in the σ 's. We define $\Delta T_t = T_t - T_t^{Ocean}$

¹⁴The general interpretation is more precisely that a_1 is the ratio $\frac{\text{initial cost of backstop}}{\text{initial cost of backstop} - \text{final cost of backstop}}$. However, for the employed value of 2 both ratios are the same so we stick with Nordhaus's interpretation.

and find

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

where $\tilde{\sigma}_{ocean}^{10 \text{ years}} = \sigma_{ocean}^{10 \text{ years}}(1 - C4) \approx 0.073 \times 0.95 \approx 0.069$. In order to obtain equation (14) stated in the main text, we have to downscale the equation to a one year time step. Assuming constant forcing and temperature feedback would result in an equation of the form $T_{t+10} = (1 - \sigma_{forc}^{dec})T_t + \Gamma$ and a downscaling of the decadal delay parameters to the one year time step as $\sigma_{forc} = 1 - (1 - \sigma_{forc}^{dec})^{\frac{1}{10}} \approx 0.032$. Linearly downscaling of $\tilde{\sigma}_{ocean}$ to the one year time step yields $\tilde{\sigma}_{ocean} \approx 0.007$. However, along the modeled paths we neither experience constant radiative forcing nor constant ocean cooling. Thus, a one to one mapping to our finer time resolution is not possible. We calibrate σ_{forc} and σ_{ocean} to obtain the best fit to the deterministic DICE model (Figure 14), but it turns out that the above approximations coincide with the calibrated values. Moreover, we extract the actual atmosphere ocean temperature difference ΔT_t from DICE-2007 and find that we can approximate and interpolate ΔT_t fairly well by the analytic function

$$\Delta T_t = \max\{0.7 + 0.02t - 0.00007t^2, 0\} . \quad (21)$$

References

- Ackerman, F., Stanton, E. A. & Bueno, R. (2010), ‘Fat tails, exponents, extreme uncertainty: Simulating catastrophe in dice’, *Ecological Economics* **69**(8), 1657–1665.
- Anthoff, D. & Tol, R. (2010), ‘Climate policy under fat-tailed risk: An application of fund’, *ESRI Working Paper* (348).
- Baker, E. & Shittu, E. (2007), ‘Uncertainty and endogenous technical change in climate policy models’, *Working Paper*.
- Basal, R., Kiku, D. & Yaron, A. (2010), ‘*’, *American Economic Review: Papers & Proceedings* **100**, 542–546.
- Basal, R. & Yaron, A. (2004), ‘Risks for the long run: A potential resolution of asset pricing puzzles’, *The Journal of Finance* **59**(4), 1481–509.
- Campbell, J. Y. (1996), ‘Understanding risk and return’, *The Journal of Political Economy* **104**(2), 298–345.
- 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.

- Ha-Duong, M. & Treich, N. (2004), ‘Risk aversion, intergenerational equity and climate change’, *Environmental and Resource Economics* **28**(2), 195–207.
- Hanemann, M. (2009), What is the economic cost of climate change?, in S. H. Schneider, A. Rosencranz, M. Mastrandrea & K. Kuntz-Duriseti, eds, ‘Climate Change Science and Policy’, Island Press, chapter 17, pp. 185–193.
- 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. & Kolstad, C. D. (1999), ‘Bayesian learning, growth, and pollution’, *Journal of Economic Dynamics and Control* **23**, 491–518.
- Leach, A. J. (2007), ‘The climate change learning curve’, *Journal of Economic Dynamics and Control* **31**, 1728–1752.
- 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.
- Richels, R. G., Manne, A. S. & Wigley, T. M. (2004), ‘Moving beyond concentrations: The challenge of limiting temperature change’, *Working Paper* .
- Traeger, C. (2009), ‘Recent developments in the intertemporal modeling of uncertainty’, *ARRE* **1**, 26185.
- Traeger, C. (2010), ‘Intertemporal risk aversion – or – wouldn’t it be nice to know whether Robinson is risk averse?’. CUDARE Working Paper No. 1102.
- Vissing-Jørgensen, A. & Attanasio, O. P. (2003), ‘Stock-market participation, intertemporal substitution, and risk-aversion’, *The American Economic Review* **93**(2), 383–391.
- Weil, P. (1990), ‘Nonexpected utility in macroeconomics’, *The Quarterly Journal of Economics* **105**(1), 29–42.
- Weitzman, M. L. (2009), ‘What is the ‘damages function’ for global warming – and what difference might it make?’, *Working Paper* .

Additional Appendix:

C Numerical Implementation and Calibration

We solve the Bellman equation (6) by function iteration. For this purpose, we approximate the value function V^* by Chebyshev polynomials. We update the coefficients by collocation at the Chebychef nodes spelled out in table 2 (rectangular grid).¹⁵ To arrive at this final node grid, we sequentially increased the number of nodes in each dimension. Figure 10 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 11 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.

Rather than modeling the carbon cycle explicitly, we assume an exogenously decreasing rate at which carbon leaves the atmosphere (reflecting that carbon sinks reduce their uptake capacity over time). We calibrate our ‘carbon model’ so that the optimal time paths of four variables are similar to those predicted by the DICE-2007 model: CO₂ concentration, temperature, abatement rate, and social cost of carbon.¹⁶ Figures 12 and 13 shows the results of calibrating the functional form stated in equation (8) to fit the depicted DICE-2007 paths. The rate decreases from the initial value $\delta_{M,0}$ to the asymptotic value $\delta_{M,\infty}$ and the rate of decline is characterized by δ_M^* . Moreover, we use an approximate interpolation and simplification in modeling the delay of atmospheric warming as described at the end of Appendix B. Figure 14 shows that the approximated values for the heat capacity delay parameter σ_{forc} and for the parameter $\tilde{\sigma}_{ocean}$ related to ocean feedback are indeed also the best fit in order to match the depicted DICE-2007 paths. Overall, these figures show that our chosen model fits the predicted paths of DICE-2007 reasonably well. Any individual parameter change that would improve the fit in one dimension, would imply a worse fit in at least one of the other variables.

For the numerical analysis under uncertainty we optimized over 5 Gaussian quadrature nodes per time step approximating the normal distribution. Figures 15 shows that increasing the number of nodes has virtually no effect on the resulting time paths. Figure 16 confirms the same finding for the ex-ante uncertainty paths for the next one and a half centuries. Beyond this point, the abatement path reflects

¹⁵For the models with joint uncertainty on damage coefficient and damage exponent we used a lower upper bound for the carbon stock and temperature intervals with the same number of nodes.

¹⁶We used the EXCEL version of the model that can be downloaded from <http://nordhaus.econ.yale.edu/DICE2007.htm> to generate the optimal time paths of DICE-2007. It generates a longer time series than depicted for example in Nordhaus (2008). Note that the EXCEL model assumes a constant savings rate. We did find an almost constant savings rate in our model as well, and the match of EXCEL DICE to the full DICE seems a close fit for the time span for which we have the comparison.

that a discrete number of paths is hitting full abatement. At these points in time the average paths exhibit kinks that vary depending on the total number of paths determining the position of the quadrature nodes. Note, however, that the social cost of carbon remains fairly insensitive to the number of quadrature nodes even beyond year 2150. Finally, Figure 17 depicts the difference between oceanic and atmospheric temperatures in DICE-2007 for a deterministic N preference run and compares it to our simple quadratic approximation stated in equation (21).

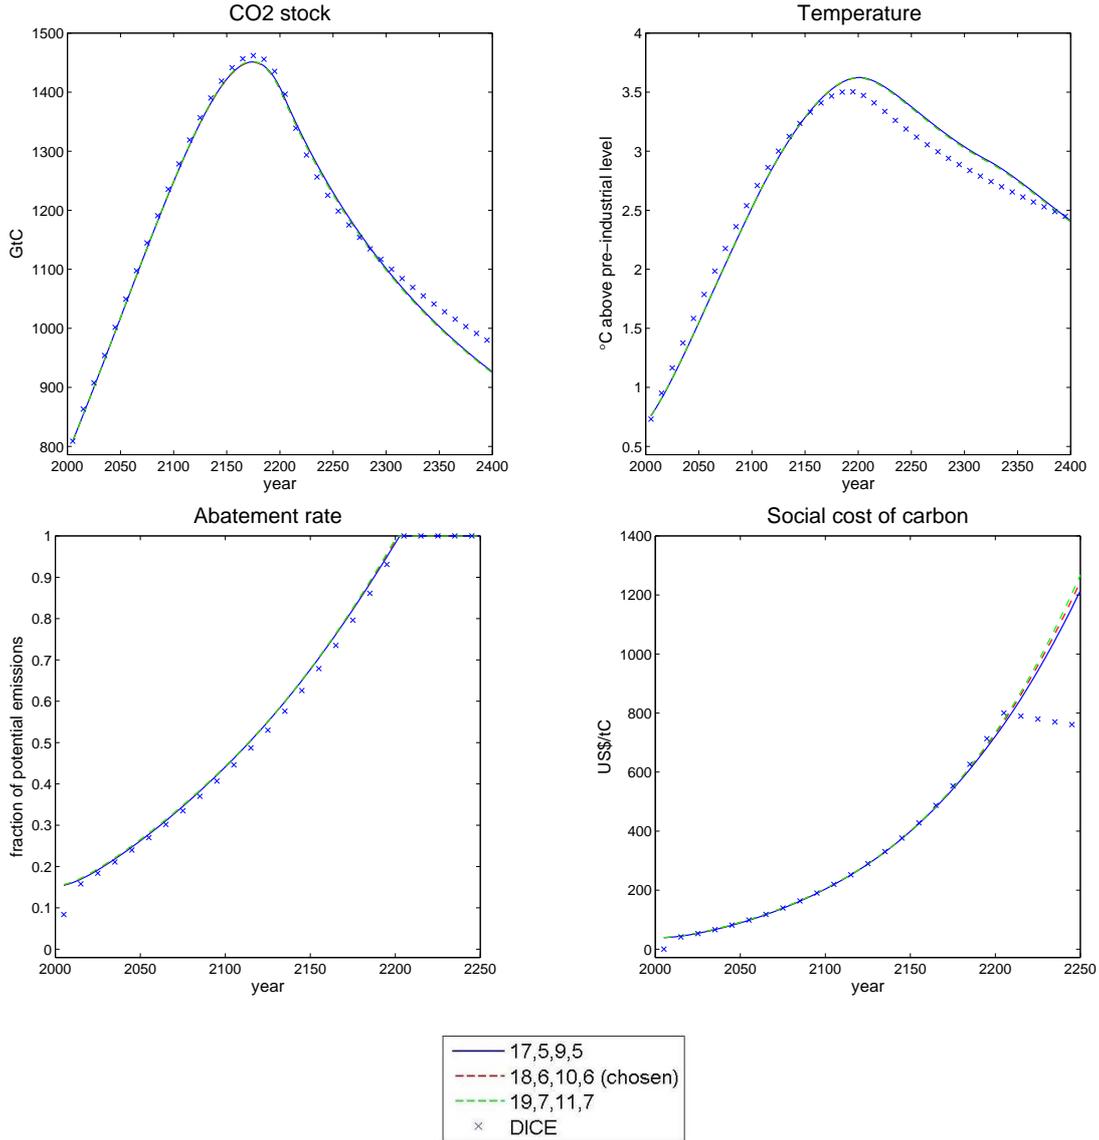


Figure 10 shows robustness of the results to variations of the number of collocation nodes.

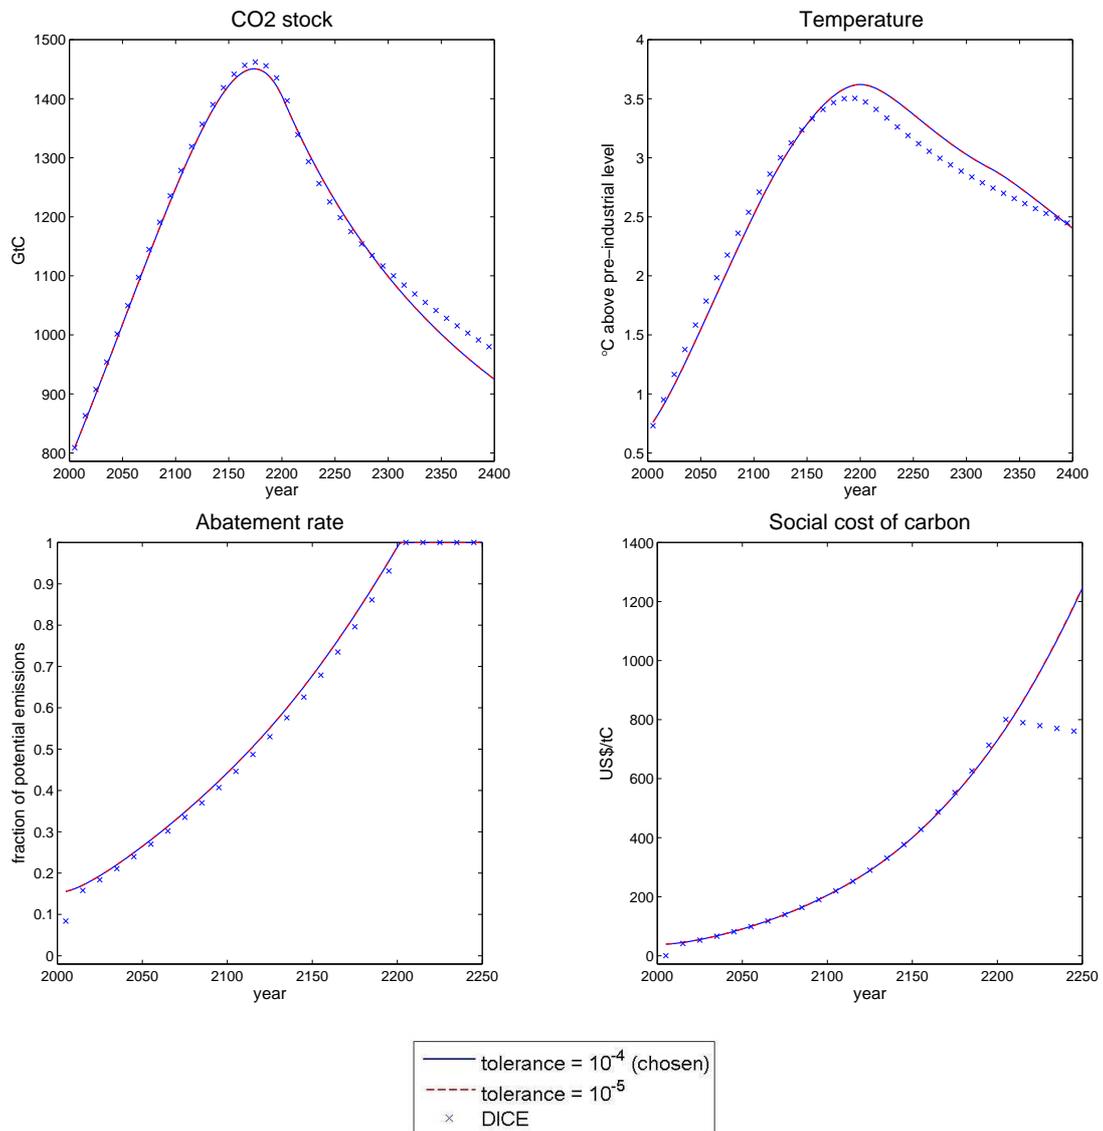


Figure 11 shows robustness of the results to a decrease in the convergence tolerance.

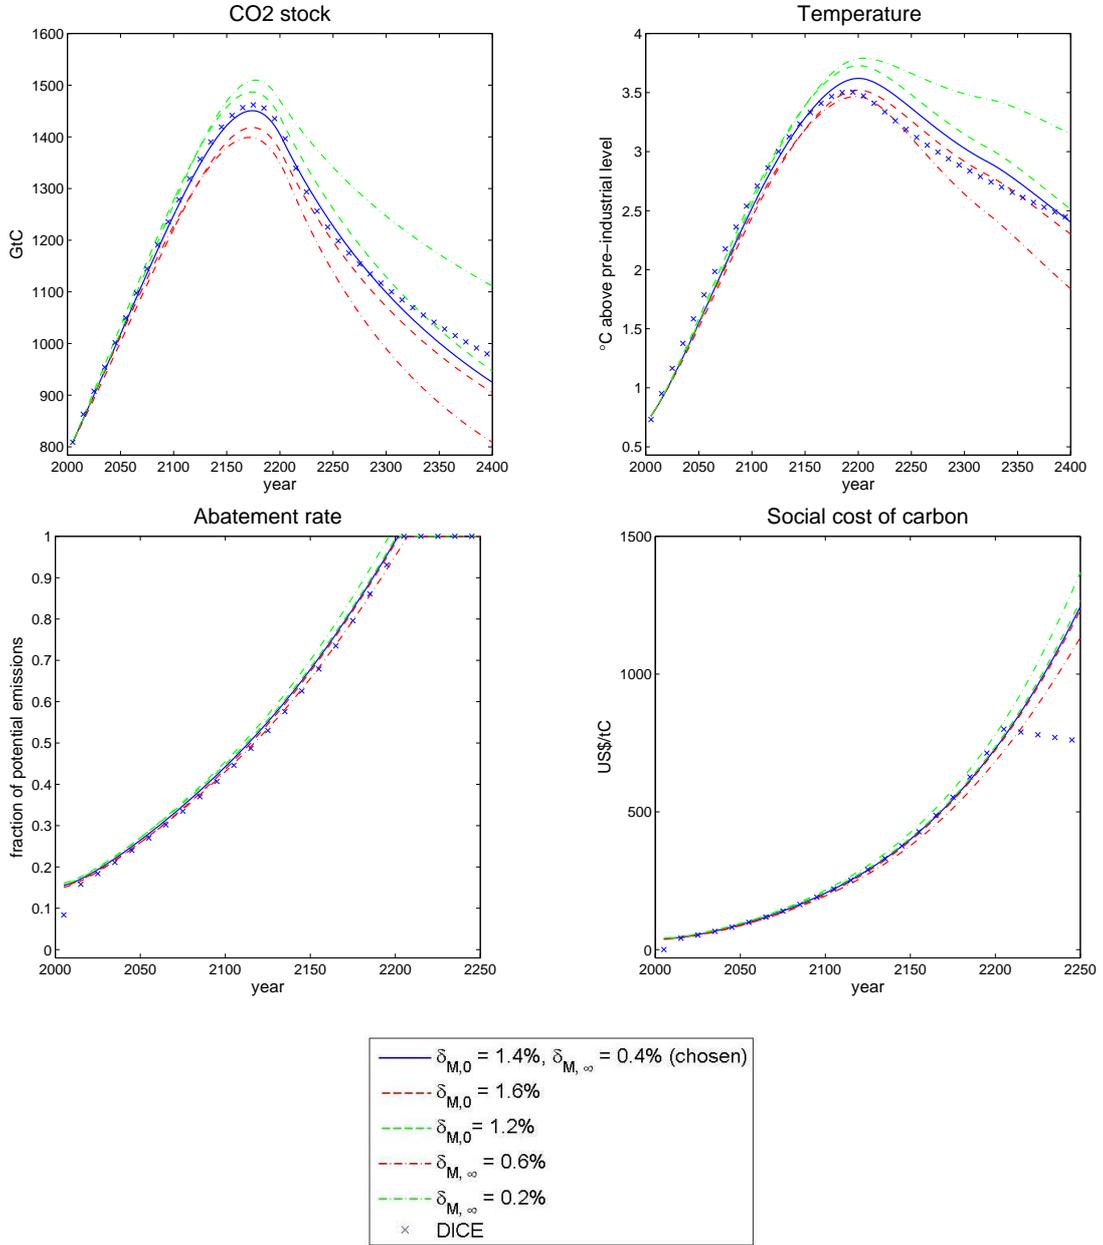


Figure 12 shows the calibration of the rate governing CO₂ removal from the atmosphere. We calibrate the initial rate $\delta_{M,0}$ and asymptotic rate $\delta_{M,\infty}$. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

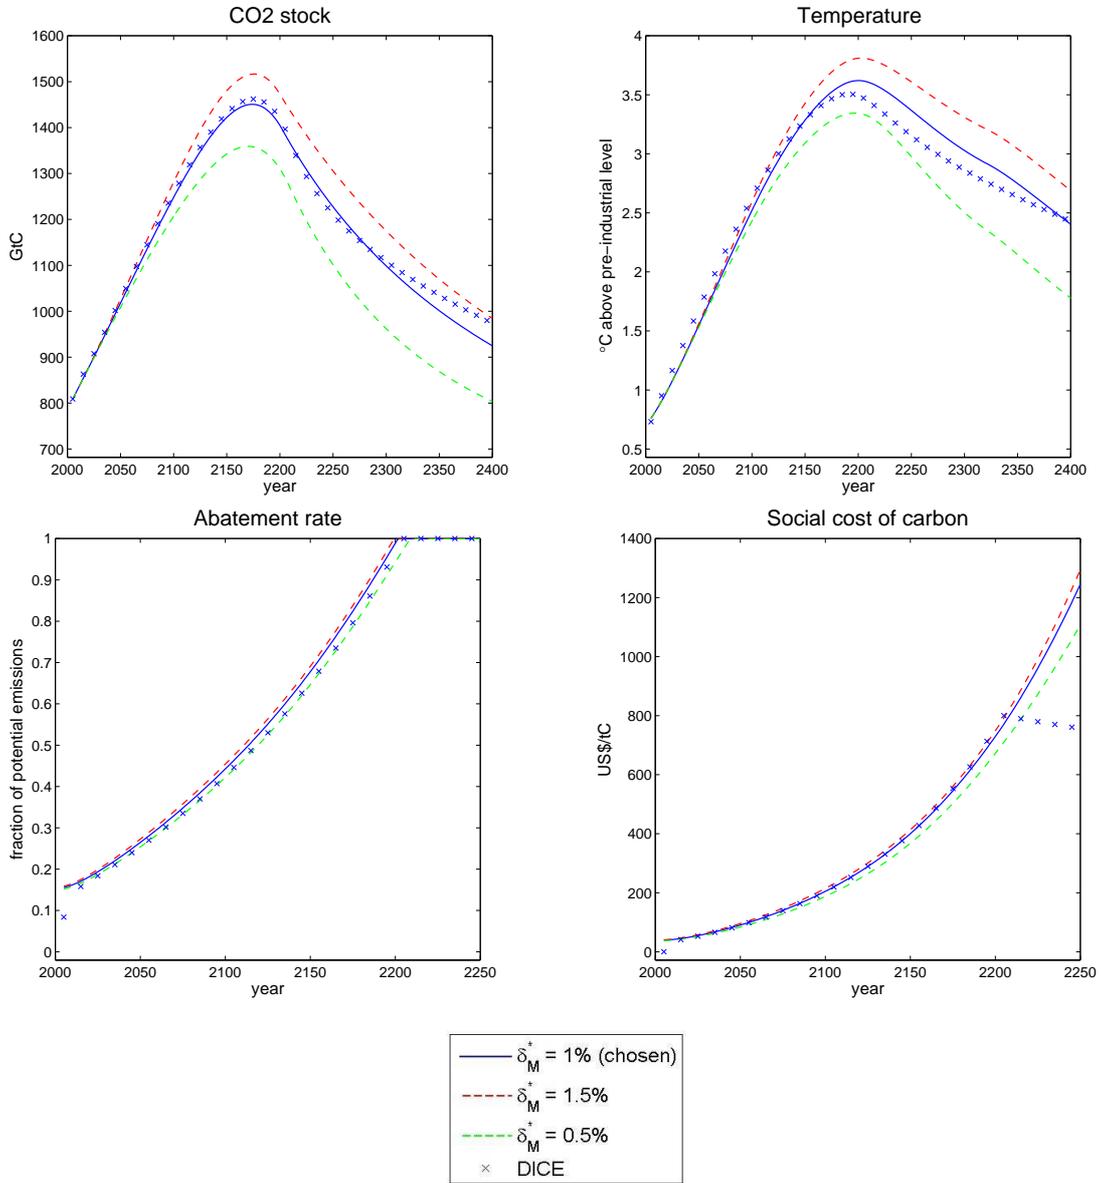


Figure 13 shows the calibration of the rate governing CO₂ removal from the atmosphere. We calibrate the parameter δ_M^* governing the speed of convergence from the initial to the asymptotic rate of CO₂ removal.

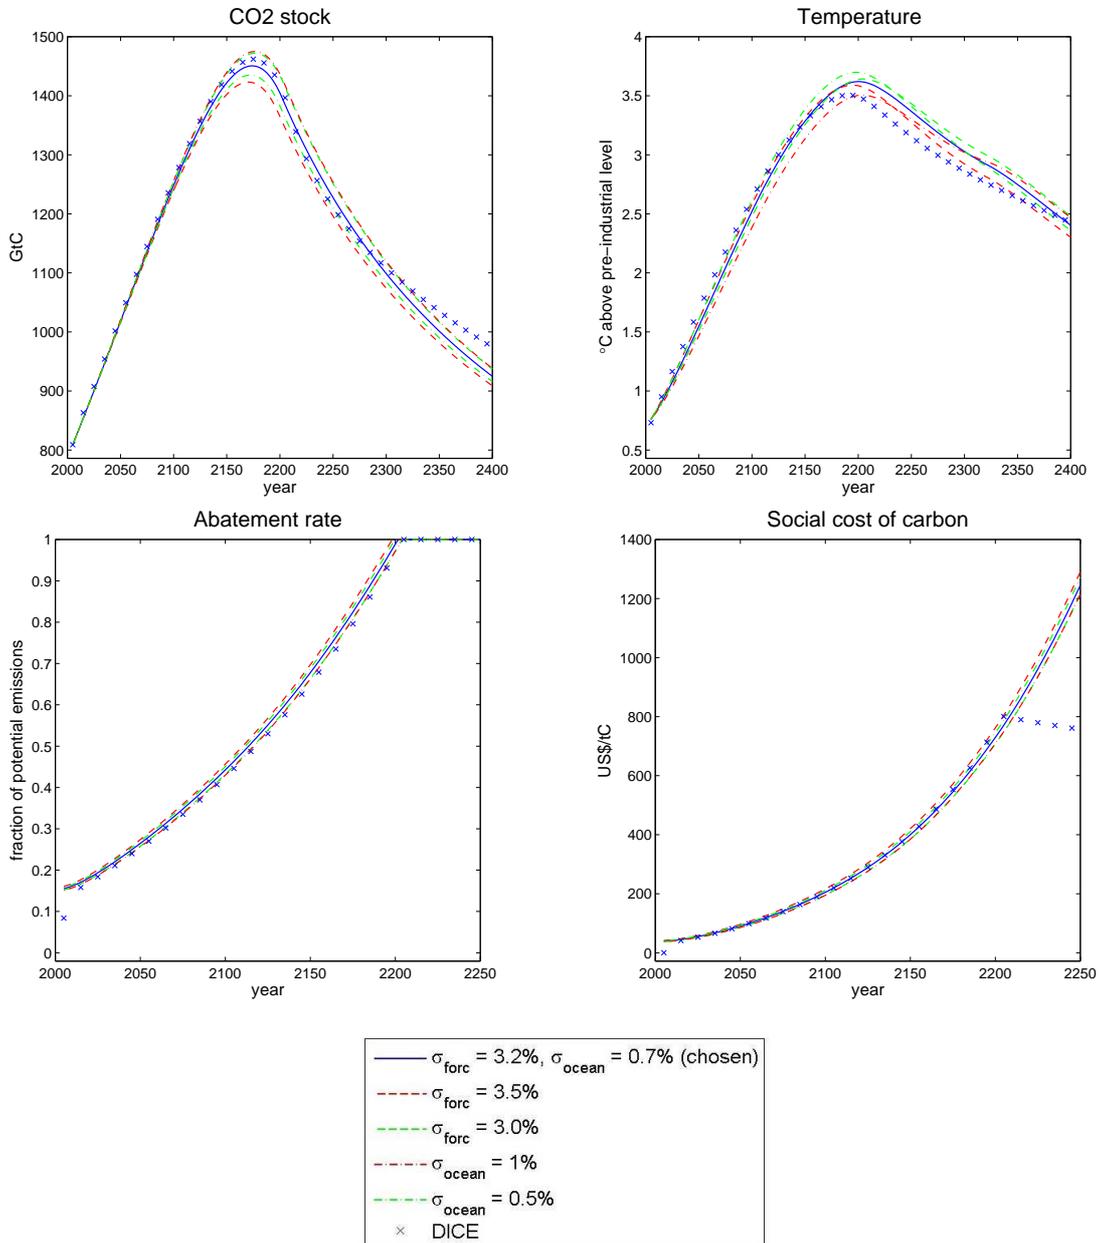


Figure 14 shows the calibration of warming delay parameter σ_{forc} and the parameter $\tilde{\sigma}_{ocean}$ connecting atmospheric and oceanic temperatures. The first line in the legend displays the parameter values chosen in our calibration. The other lines show the value of the parameter that was changed with respect to our chosen calibration.

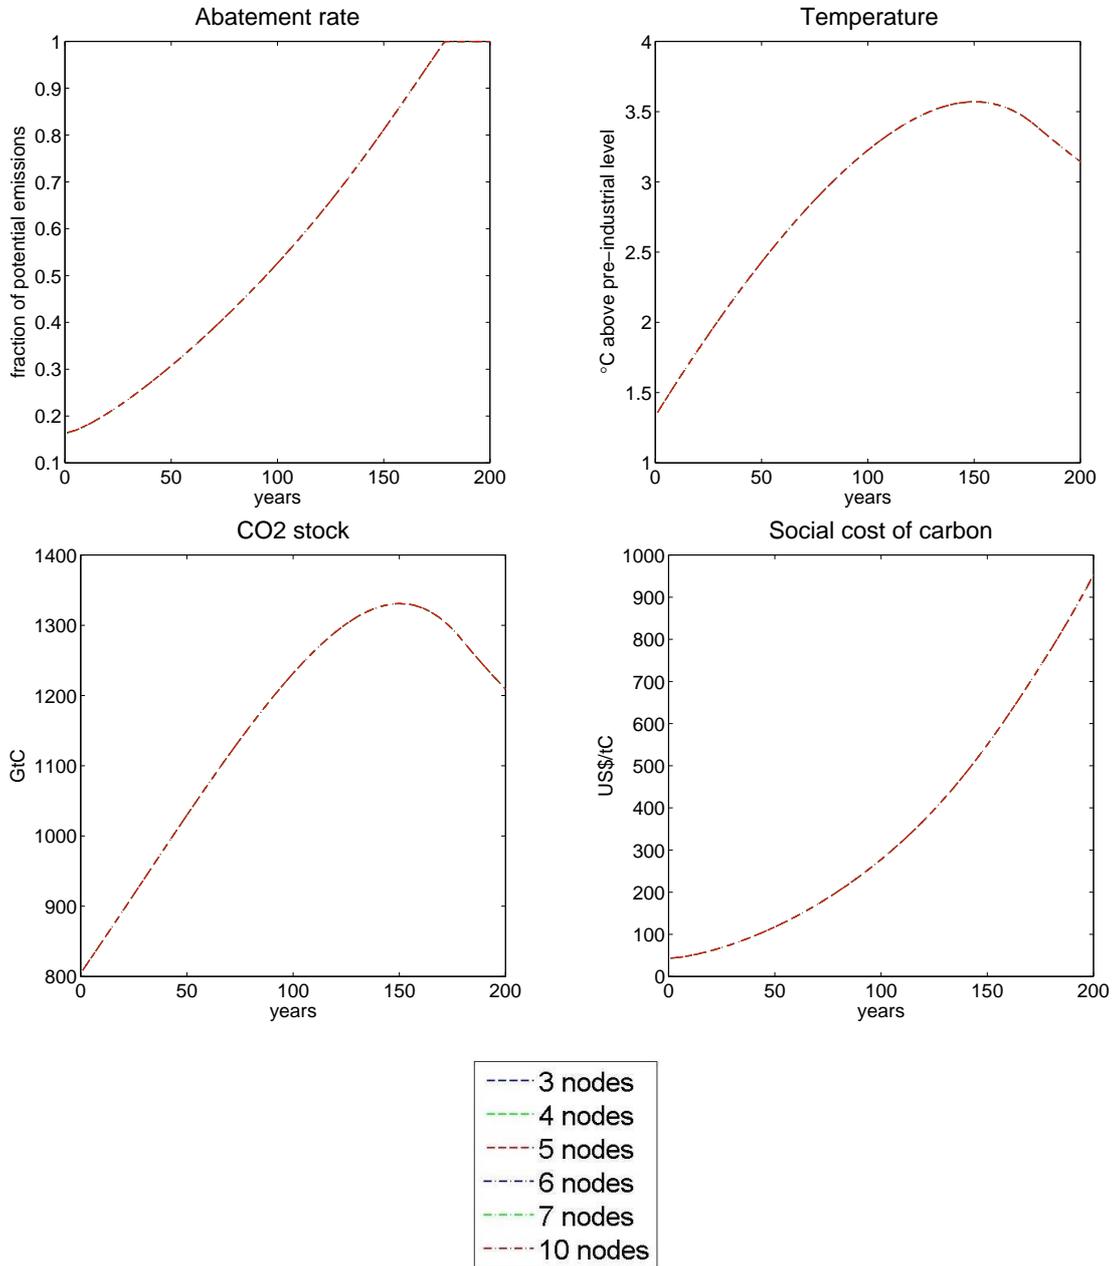


Figure 15 shows the robustness of our stochastic dynamic programming results to variations in the number of Gauss-Legendre quadrature nodes.

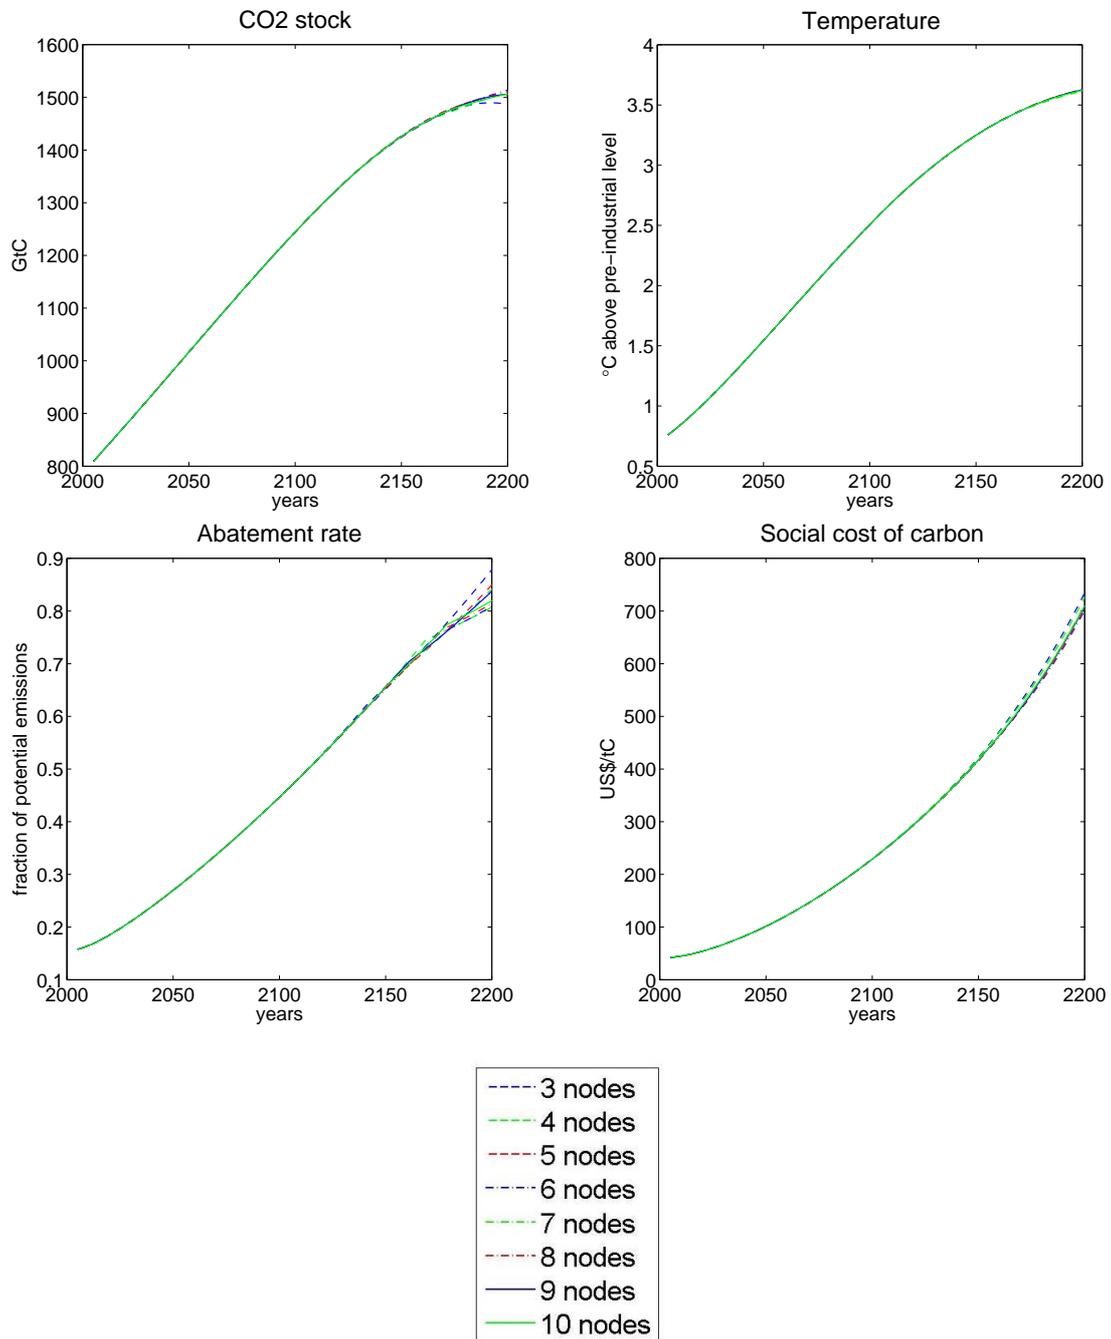


Figure 16 shows the effect of increasing the number of quadrature nodes on the results of our ex ante uncertainty simulation.

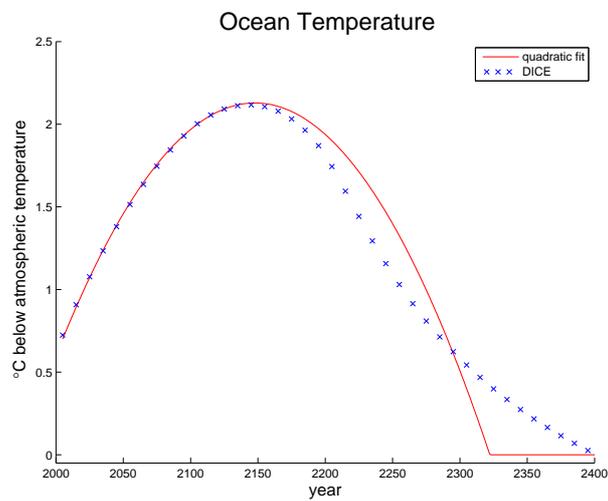


Figure 17 compares our simple quadratic approximation of the temperature difference between oceans and atmosphere to the actual difference resulting from the DICE-2007 model.

Table 1 Parameters of the model

Economic Parameters		
η	$\frac{2}{3}, 2$	intertemporal consumption smoothing preference
RRA	0, 2, 9.5, 50	coefficient of relative Arrow-Pratt risk aversion
b_1	0.00284	damage coefficient; for uncertain scenario normally distributed with standard deviation 0.0013 (low) and 0.0025 (high)
b_2	2	damage exponent; for uncertain scenario normally distributed with standard deviation 0.35 (low) and 0.5 (high)
δ_u	1.5%	pure rate of time preference
L_0	6514	in millions, population in 2005
L_∞	8600	in millions, asymptotic population
g_L^*	0.035	rate of convergence to asymptotic population
K_0	137	in trillion 2005-USD, initial global capital stock
δ_K	10%	depreciation rate of capital
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
δ_A	0.1%	rate of decline of productivity growth rate
σ_0	0.1342	CO ₂ emission per unit of GDP in 2005
$g_{\sigma,0}$	-0.73%	initial rate of decarbonization
δ_σ	0.3%	rate of decline of the rate of decarbonization
a_0	1.17	cost of backstop 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, preindustrial stock of CO ₂ in the atmosphere
M_0	808.9	in GtC, stock of atmospheric CO ₂ in 2005
$\delta_{M,0}$	1.4%	initial rate of decay of CO ₂ in atmosphere
$\delta_{M,\infty}$	0.4%	asymptotic rate of decay of CO ₂ in atmosphere
δ_M^*	1%	rate of convergence to asymptotic decay rate of CO ₂
B_0	1.1	in GtC, initial CO ₂ emissions from LUCF
g_B	-1%	growth rate of CO ₂ emission from LUCF
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
$\tilde{\sigma}_{ocean}$	0.7%	parameter governing oceanic temperature feedback

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			