

University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

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

This Version: June 2010 , First Version: January 2010

!Print preferably in COLOR!

Abstract: We analyze the impact of uncertainty on optimal mitigation policies derived from the integrated assessment of climate change. For this purpose, we construct a close relative of the DICE model in a recursive dynamic programming framework. First, our framework can capture persistent uncertainty and we compare it to the simpler and more frequently employed analysis of ex-ante uncertainty. Second, our framework makes it possible to disentangle effects deriving from risk, from risk aversion, and from a decision maker's aversion to intertemporal substitution. We analyze uncertainty over climate sensitivity as well as over damages.

JEL Codes: Q54, Q00, D90, C63

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

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 about climate change taking place today and, in particular, in the future. The details of these changes and their socio-economic consequences remain uncertain at least for several more decades. The vast majority of integrated assessment models in climate change does not take proper account of the risks and uncertainties involved. In particular, when these models considers uncertainty at all, most of them tend to capture uncertainty only as ex-ante risk over model parameters resolving already before the actual climate policy starts. Moreover, these models only capture risk aversion to the extent that it is produced by the desire to smooth consumption over time.

We replicate a slightly simplified version of Nordhaus's (2008) DICE model as a recursive dynamic programming model. Our model features persistent uncertainty in an annual time step under an infinite planning horizon. We not only derive optimal expected trajectories of the climate economy and confidence bounds, but we also obtain the optimal control rules as functions of the stock variables. For example, they give the optimal carbon tax at a given point in the future conditional on the capital and carbon stocks (which depend on the earlier shock realizations). We focus on uncertainty about climate sensitivity and uncertainty over the damages caused by a given temperature increase. The climate sensitivity parameter characterizes the temperature response to the radiative forcing caused by a doubling in atmospheric CO₂ concentrations with respect to the preindustrial level. It is one of the key unknowns in modeling global warming because it depends on a number of feedback processes. With respect to damages, we model two different ways how uncertainty enters the equation. Both differ significantly in their impact on the optimal mitigation policy. A collection of arguments why we are convinced that the DICE damage function in DICE is at best a rather uncertain best guess are for example found in Hanemann (2009).

We employ Epstein & Zin's (1989) recursive utility approach in order to distinguish between a decision maker's propensity to smooth consumption over time and his Arrow-Pratt risk aversion. First, this approach makes it possible 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.¹ Figure 1 on page 10 can serve as a preview and motivation comparing abatement rate and social cost of carbon for certain and uncertain runs using the original DICE parameterization (blue) versus preference parameters chosen according to Vissing-Jørgensen & Attanasio's (2003) best estimate of actual preferences disentangling risk aversion from the desire to smooth consumption over time (green). The solid lines represent certainty while the dashed lines introduce

¹Also from a normative rather than an observational point of view we are not aware of a convincing argument that these a priori different dimensions of preference should coincide. However, it would be easy to argue for either of the two aversions being larger.

uncertainty over climate sensitivity.

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 distribution and runs DICE for each realization. While this way of modeling uncertainty is an interesting first approximation, it certainly does not describe the type of uncertainty surrounding climate change. In such a procedure uncertainty only exists right before the model is run and before a policy has to be passed. Once the model is initiated all uncertainty is resolved and taxes are set optimally under full certainty. The procedure only gives an estimate for an ex-ante prediction of a tax, which, in general, will never be implemented because the underlying world has no uncertainty left once the policy has to be passed. Ackerman, Stanton & Bueno (2010) modify Nordhaus's (2008) approach. 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, but comparing to Nordhaus (2008) they not only increase uncertainty but also expected values. Unfortunately the model does not translate the results of the simulations into changes of the optimal climate policy like for example the optimal carbon tax.

Closest to our approach are two papers by Kelly & Kolstad (1999) and Ha-Duong & Treich (2004). In their seminal paper Kelly & Kolstad (1999) implement a recursive version of DICE to analyze learning about climate sensitivity. Their setting, however, cannot disentangle risk aversion from intertemporal substitutability. Our analysis focusses on separating the impacts of risk, risk aversion, and intertemporal substitutability as well as giving best estimates for the change in the optimal carbon tax under a better preference representation than possible in the standard model. None of these aspects is analyzed by Kelly & Kolstad (1999). Moreover, they limit the analysis to uncertainty over climate sensitivity while we also examine damage uncertainty in some detail. 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 blown infinite horizon integrated assessment model and add a quite realistic description of uncertainty. We derive the actual magnitude of the effects of risk, risk aversion, and aversion to intertemporal substitution affect the social cost of carbon and compare results under the preference parameters used by Nordhaus (2008) to those based on actual estimates in the asset pricing literature using the disentangled approach.

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

Moreover, we derive the optimal control rules. Section 2 introduces the recursive utility specification and explains the disentanglement of risk aversion from intertemporal substitutability. Section 3 introduces the climate enriched economy. Section 4 presents the results and section 5 concludes. Additional information on our approach is summarized in an appendix.

2 Welfare and Bellman equation

2.1 The Original Welfare Specification

In this section we introduce our welfare function and the dynamic programming equation. The main feature distinguishing our setting from other integrated assessment models is the ability to disentangle effects of risk and risk attitude from effects driven by the desire to smooth consumption over time. For this purpose we employ a variant of the generalized isoelastic model introduced by Epstein & Zin (1989) and Weil (1990) disentangling Arrow Pratt risk aversion from the aversion to intertemporal substitution. Such a disentanglement is not possible using the standard intertemporally additive expected utility approach of the form $U = E \sum_t \exp[-\delta_u t] u(x_t)$ evaluating scenarios by aggregating instantaneous welfare linearly over time and over risk. In such a model, the instantaneous utility function u is used as an aggregator in both dimensions, time and risk. Therefore, the concavity of u simultaneously captures aversion to intertemporal substitution as well as to risk aversion. A priori, however, these two characteristics of preference are very different. Empirical estimates disentangling the two dimension and surveys of such include Campbell (1996), Giuliano & Turnovsky (2003), and Vissing-Jørgensen & Attanasio (2003). (Traeger 2007) traces back the difference to simple intuitive axioms concerning lottery choices. The evidence suggests that people tend to be more averse to substituting consumption into a risk state than into the certain future, implying higher Arrow Pratt risk aversion than 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 employ a recursive utility. In consequence, we cannot use the standard optimal control framework employed for solving most integrated assessment models. The Bellman equation in our setting writes as

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

The value function V represents the value of an optimal path given the following state variables: Time t , capital K_t , and stock of carbon in the atmosphere M_t . Utility within a period corresponds to the first term on the right hand side of the dynamic

programming equation (1) and depends on total global consumption C_t per capita and the population L_t . The parameter η captures the desire to smooth consumption over time or aversion to intertemporal substitution. It is the inverse of the intertemporal elasticity of substitution. The parameter RRA depicts the Arrow Pratt measure of relative risk aversion. The standard model implicitly assumes $\eta = \text{RRA}$. The factor $\exp[-\delta_u]$ discounts future welfare (one period ahead) with the rate of pure time preference δ_u , which we assume to be 1.5% in accordance to Nordhaus (2008) throughout the paper. The decision maker maximizes over the control variables C_t and the emission control rate μ_t .³ Thus, 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 of the path in the next period given the new state variables (which are altered according to the consumption and abatement decision).

For a detailed analysis of the interpretation of the parameters α and ρ I refer to Epstein & Zin (1989) and to Traeger (2007) who also derives the particular representation used above that is additive in the time step. As our baseline we employ a scenario based on Nordhaus's (2008) preference parameters $\eta = \text{RRA} = 2$. We contrast it with a best guess estimate of disentangled preferences by Vissing-Jørgensen & Attanasio (2003) who build on Campbell's (1996) approach estimating the Epstein & Zin (1989) preferences by means of a log-linearized Euler equations in the asset pricing context yielding $\eta = \frac{2}{3}$ and $\text{RRA} = 9.5$. For both scenarios we also vary the risk aversion coefficient from risk neutrality $\text{RRA} = 0$ up to extreme risk aversion $\text{RRA} = 50$.

2.2 A Conveniently Modified Bellman Equation

Our integrated climate economy is constructed as a close match to Nordhaus (2008) and features exogenous technological progress and population growth. We improve the performance of the recursive numerical model significantly by expressing the relevant variables in effective labor terms and rewriting the Bellman equation accordingly. Exogenous technological progress is characterized by the variable

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

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

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

³In the numerical implementation of the model it turns out useful to maximize over the abatement cost Λ_t , which is a strictly monotonic transformation of μ_t , in order to simplify the model constraints.

L_0 denotes the initial and L_∞ the asymptotic population. The parameter δ_L^* characterizes the convergence from initial to asymptotic convergence.⁴ Figure 8 in the appendix summarizes the time behavior of the exogenous time paths, all of which we took over from 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 definition a couple of manipulations transform equation (1) into

$$\frac{V(k_t A_t L_t, M_t, t)}{A_t^\rho L_t} = \max_{c_t, \mu_t} \frac{c_t^\rho}{\rho} + \frac{\exp[-\delta_u + g_{A,t} \rho + g_{L,t}]}{\rho} \left(\mathbb{E} \left[\rho \frac{V(k_{t+1} A_{t+1} L_{t+1}, M_{t+1}, t+1)}{A_{t+1}^\rho L_{t+1}} \right]^\alpha \right)^{\frac{\rho}{\alpha}} \quad (2)$$

with the utility discount factor $\beta = \exp[-\delta_u]$ and consumption per effective labor $c_t = \frac{C_t}{A_t L_t}$.

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

$$\tau = 1 - \exp[-\zeta t] \in [0, 1] \quad (3)$$

and define

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

where A_t and L_t follow

$$A_t = A_0 \exp \left[g_{A,0} \frac{1 - \exp[-\delta_A t]}{\delta_A} \right] \text{ and } L_t = L_0 + (L_\infty - L_0)(1 - \exp[-\delta_L^* t]) .$$

Here and in the following we have denoted under slight abuse a variables x_τ to represent the variable $x_{\tau(t)}$ under the transformation 3. Then, we can rewrite the fixed point equation (2) in terms of V^* as

$$V^*(k_\tau, M_\tau, \tau) = \max_{c_\tau, \mu_\tau} \frac{c_\tau^\rho}{\rho} + \frac{\exp[-\delta_u + g_{A,\tau} \rho + g_{L,\tau}]}{\rho} \left(\mathbb{E} [\rho V^*(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, \tau + \Delta\tau)]^\alpha \right)^{\frac{\rho}{\alpha}}$$

⁴More precisely the discrete time growth rate would be $\ln \left[1 + \frac{1 - \exp[\delta_L^*]}{(1 + \frac{L_0}{L_\infty}) \exp[\delta_L^* t] - 1} \right]$. We use the above approximation.

⁵Note that the time transformation also concentrates the (optimally spread) nodes at which we evaluate our Chebychev polynomials (approximating the value function) on the close future in real time, where most of the exogenously driven changes take place.

Note that the time step is now in artificial time. Keeping 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 $\tau + \Delta\tau = 1 - [1 - \tau] \exp[-\zeta]$.

3 The Climate Economy

3.1 The model under certainty

The decision maker maximizes his value functions under the constraints of the following stylized model of a climate enriched economy. The model is largely a reproduction of Nordhaus (2008) DICE-2007 model except for two simplifications. Because state variables are computationally intensive in a recursive dynamic programming model we neither make temperature a state variable nor do we model CO2 concentrations in the oceans explicitly. The first simplification does not permit us to capture the delay between a radiative forcing increase and a temperature increase (caused by feedback processes). The second simplification replaces the simple carbon cycle model in DICE by an exogenously falling decay rate, which mimics that other reservoirs reduce their take-up rate as they fill up. We calibrate our exogenous change in the decay rate to match the CO2 time path of DICE-2007 under certainty with Nordhaus's parameter specifications. The only other modification to the DICE model is that we focus solely on CO2 emissions and chose not to include the exogenous forcing from other greenhouse gases that DICE includes as a quite simple exogenous path that starts out cooling the system today and warming the system additionally in the future.⁶ In the following we discuss the underlying equations. While some of the equations read different in the recursive setting and in DICE, they are mathematically equivalent except for the fact that we use a yearly time step. All parameters are characterized and quantified in table 3.1 on page 24.

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. Instead of trying to model the full carbon cycle, which would be very costly in terms of stock variables, we assume an exponential decay of CO₂ in the atmosphere at rate $\delta_{M,t}$ which is exogenously reduced over time to replicate the carbon stock of DICE-2007 featuring a carbon cycle (see figure 7 in the appendix)

$$\begin{aligned} M_{\tau+\Delta\tau} &= M_\tau (1 - \delta_{M,\tau}) + E_\tau \quad \text{with} \\ \delta_{M,t} &= \delta_{M,\infty} + (\delta_{M,0} - \delta_{M,\infty}) \exp[-\delta_M^* t] . \end{aligned} \tag{4}$$

⁶The effect of the exogenous forcing path is very minor in DICE, but a rigorous treatment of the other greenhouse gases might likely require a somewhat deeper analysis.

The variable E_t characterizes overall yearly CO₂ emissions. Emission are composed of industrial emission (first term) and emissions from land use change an forestry B_τ

$$E_\tau = (1 - \mu_\tau) \sigma_\tau A_\tau L_\tau k_\tau^\kappa + B_\tau , \quad (5)$$

where σ_t characterizes the baseline decarbonization of production. It follows the exogenous time path

$$\sigma_t = \sigma_0 \exp[g_{\sigma,t} t] , \quad \text{with} \quad g_{\sigma,t} = \exp[-\delta_\sigma t] ,$$

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

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

Together we can rewrite the equation for emissions (5) as

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

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

$$y_\tau = \frac{1 - \Lambda(\mu_\tau)}{1 + D(T_\tau)} k_\tau^\kappa = \frac{1 - a_1 \mu_\tau^{a_2}}{(1 + b_1 T_\tau + b_2 T_\tau^{b_3})} k_\tau^\kappa \quad (6)$$

where

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

characterizes abatement costs as percent of GDP depending on the emission control rate $\mu_t \in [0, 1]$. The cost function coefficient is time dependent and is given by

$$\Psi_t = \frac{\sigma_t}{a_2} a_0 (1 - a_1 (1 - \exp[g_\Psi t])) \quad (7)$$

with a_0 denoting the initial cost of backstop (i.e. in 2005), a_1 denoting the ratio of initial over final backstop,⁷ and a_2 is the cost exponent. The rate g_Ψ describes

⁷The 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.

the convergence from initial to final cost of backstop. In artificial time equation (7) translates into

$$\Psi_\tau = \frac{\sigma_0 a_0}{a_2} \exp \left[g_\sigma \frac{1 - (1 - \tau)^{\frac{\delta_\gamma}{\zeta}}}{\delta_\gamma} \right] \left(1 - a_1 \left(1 - [1 - \tau]^{-\frac{g_\Psi}{\zeta}} \right) \right) \quad (8)$$

Climate damage as percent of GDP depending on the temperature difference T_t of current with respect to preindustrial temperatures are characterized by

$$D(T_\tau) = b_1 T_\tau + b_2 T_\tau^{b_3}$$

where Nordhaus (2008) estimates $b_1 = 0$ and $b_3 = 2$ implying a quadratic damage function. As we do not capture delay our temperatures is an immediate response to the radiative forcing caused by the stock of CO₂ in the atmosphere

$$T_\tau = s_\tau \frac{\ln \frac{M_\tau}{M_{preind}}}{\ln 2},$$

where s_τ denotes climate sensitivity, i.e. the temperature response to a doubling of CO₂ in the atmosphere with respect to preindustrial concentrations.

3.2 Uncertainty

We introduce two sources of uncertainty into the model. First, we capture uncertainty about climate sensitivity, i.e. the temperature response to an increase in the radiative forcing cause by the increase in CO₂ concentration in the atmosphere. We assume climate sensitivity to be lognormal distribution with $\mu = 1$ and $\sigma = .5$ implying an expected value of 3.08 and a standard deviation of approximately 2.7. This distribution is more risky (in the Rothschild & Stiglitz (1970) sense) than the assumed uncertainty in the according runs by Nordhaus (2008). It is based on visually “calibrating” a lognormal to the set of pdfs suggested in IPCC (2001, Fig TS.25 p 65).⁸ Second, we analyze uncertainty about the damages. To these ends we consider a normal distribution over the damage coefficient b_2 as well as a normal over the damage exponent b_3 . For b_2 we follow Nordhaus’s (2008) assumption of a standard deviation of 0.0025% (with the mean being 0.00284). Nordhaus does not consider uncertainty over the exponent and we use a standard deviation of 0.35% (around the mean 2) that gives roughly the same damages as the b_2 uncertainty for the bounds +/- standard deviation at a 3°C temperature increase.⁹

⁸Nordhaus (2008) assumes a normal with a standard deviation 1.11, which in relation to IPCC (2001, Fig TS.25 p 65) is outperformed by our visual best fit.

⁹Precisely, we calibrate closer to the upper than to the lower one sigma deviation for a 3°C increase. The damage for $b_3 = 2 - 0.35$ results in a 1.74% loss of GDP as opposed to 1.39% in the Nordhaus (2008) based b_2 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_3 uncertainty scenario.

In our uncertain scenarios, the value function is calculated optimizing over 8 quadrature nodes per time step representing a log/normal distribution. Thus, the decision maker's control rules are based on a decision tree for which every subtree every year divides into 8 further subtrees (and 64 with joint uncertainty over two parameters). The most precise way to represent the decisions under uncertainty is by depicting the control rules. However, time paths of abatement and the social cost of carbon are insightful representations that we do not forgo. For that purpose we will generally depict the particular time path that comes about when the decision maker decides under uncertainty, but nature always happens to draw the expected value of the distribution. While we use that particular path to compare different scenarios, we will also examine closely how this path relates to the median path and the expected path in the sense of ex-ante uncertainty a la Nordhaus (2008). The latter draws the climate sensitivity parameter randomly before running the model and then runs the model for each draw as if climate sensitivity would be certain. Taking expectations over these runs is what we label ex-ante uncertainty. Such a model of ex-ante uncertainty corresponds to a situation where the decision maker is only uncertain about which climate system he confronts *before* he starts making decisions. As soon as he starts to control the system in $t = 0$ all uncertainty has resolved.

3.3 Remarks on Delay, Calibration, and Effects on the Carbon Tax

As we mentioned above we do not capture temperature delay. Such a delay pushes the negative effects of emissions into the future and reduces the net present value of their damage. Thus, we slightly overestimate the social cost of carbon. However, we also calibrate our carbon decay rate to mimic the DICE model in the certain baseline run. This calibration does more than simply matching the carbon cycle. Without temperature delay our model would tend to have slightly lower emissions than DICE while temperature is on the rise. By calibrating the time path of the CO2 stock to that of DICE, we effectively choose a slightly higher decay rate in the early decades than a real carbon cycle model would bring about. Such a slightly higher rate of decay reduces the social cost of carbon, balancing off some of the effect from not modeling temperature delay.

4 Results

Figure 1 presents the optimal abatement rate as well as the social cost of carbon (SCC) or optimal carbon tax for the parameter choice of Nordhaus (2008) (dark blue) as well as for the assessment using disentangled preferences based on the parameter estimates of Vissing-Jørgensen & Attanasio (2003) (light green). The solid lines represent assessment under certainty. The dashed lines introduce uncertainty over

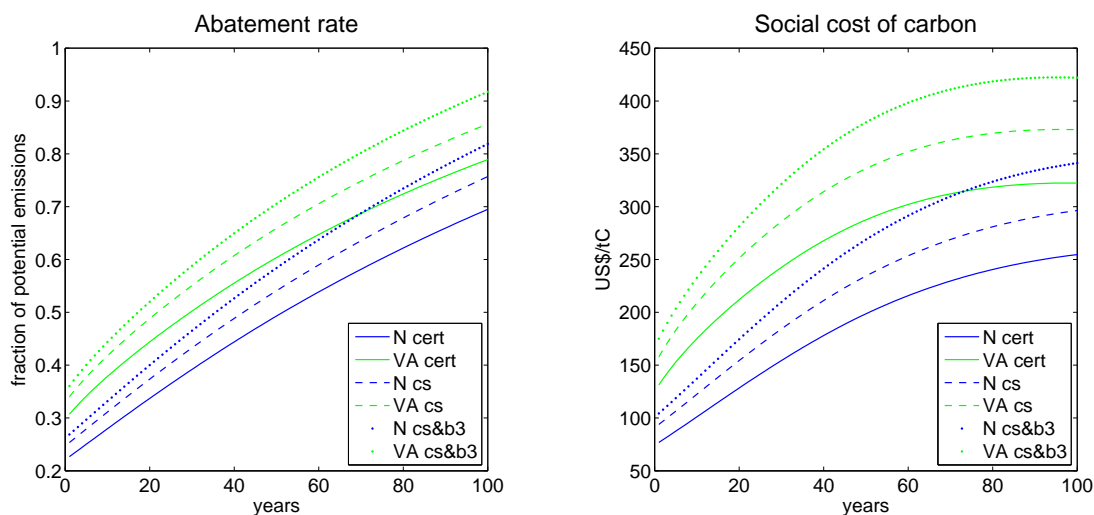


Figure 1 compares the optimal abatement rate and CO₂ tax based on the standard model with parameters $\eta = RRA = 2$ as in DICE-2007 (N, dark blue lines) with an assessment based on the disentangling model with parameter estimates $\eta = \frac{2}{3}$ and $RRA = 2$ taken from Vissing-Jørgensen & Attanasio (2003) (VA, light green lines). “cert” denotes assessment under certainty, “cs” introduces uncertainty over climate sensitivity (dashed lines), and “cs&b3” introduces (independent) uncertainty over climate sensitivity and the damage exponent (dotted lines).

climate sensitivity, and the dotted lines give optimal abatement and SSC for the scenario with both, uncertainty over climate sensitivity and over damages (precisely over the damage exponent b_3). For the uncertain scenarios the figure depicts the path where the decision maker chooses under uncertainty and nature happens to draw expected values.

Both, the abatement rate and the optimal carbon tax, are higher for in the case of disentangled preferences. In numbers, the additional SCC over the next decade is about \$60 to \$100 higher under the disentangled preference scenarios than with the Nordhaus preferences. In relative terms the optimal tax is about 70% higher at the beginning of the century and about 25% higher at the end of the century. Over the next 10 to 100 years uncertainty over climate sensitivity increases the optimal carbon tax \$20 to \$40 with Nordhaus preferences and \$35 to \$50 with disentangled preferences. Including as well uncertainty over the damage exponent adds once more a similar amount to the SCC. A first conclusion from Figure 1 is that, first, uncertainty yields a significant effect on abatement and optimal taxes and that, second, a distinction between risk preferences and the propensity to smooth consumption over time results in a significant change of the optimal policy. In what follows we examine more closely the underlying cause of the differences observable in Figure 1. In particular, we will analyze whether the risk effects are driven mostly by risk or mostly by risk aversion, and whether the risk aversion or the propensity to smooth consumption over time are more important determinants of the optimal carbon tax.

Figure 2 examines the dependence of the abatement rate and the SCC on the




Figure 2 compares the optimal abatement rate and optimal carbon tax for different degrees of risk aversion (RRA) in the the scenario with $\eta = 2$ as in DICE-2007 (upper graphs) and for $\eta = \frac{2}{3}$ as suggested by Vissing-Jørgensen & Attanasio (2003) in the disentangled approach (lower graphs). Except for the red lines all scenarios feature uncertainty over climate sensitivity.

degrees of risk aversion. Thanks to the recursive preferences we can vary risk aversion (RRA) while keeping aversion to intertemporal substitution (η) constant. The above figures fix $\eta = 2$ as in DICE-2007, while the lower figures fix $\eta = \frac{2}{3}$ as suggested by Vissing-Jørgensen & Attanasio (2003). The difference between the certain paths (red) and the optimal paths under uncertainty shows that risk over climate sensitivity has a major effect on the optimal policies under any assumption on risk aversion. The optimal policy hardly differs between the cases of risk neutrality, of RRA = 2 as in DICE-2007, and of RRA = 9.5 as suggested by the disentangled estimates. Only an extreme degree of risk aversion of RRA = 50 has a non-negligible effect on the policies. However, even as an upper bound on risk aversion RRA = 50 very high. Thus, we conclude from Figure 2 that risk aversion itself is not the driver of the

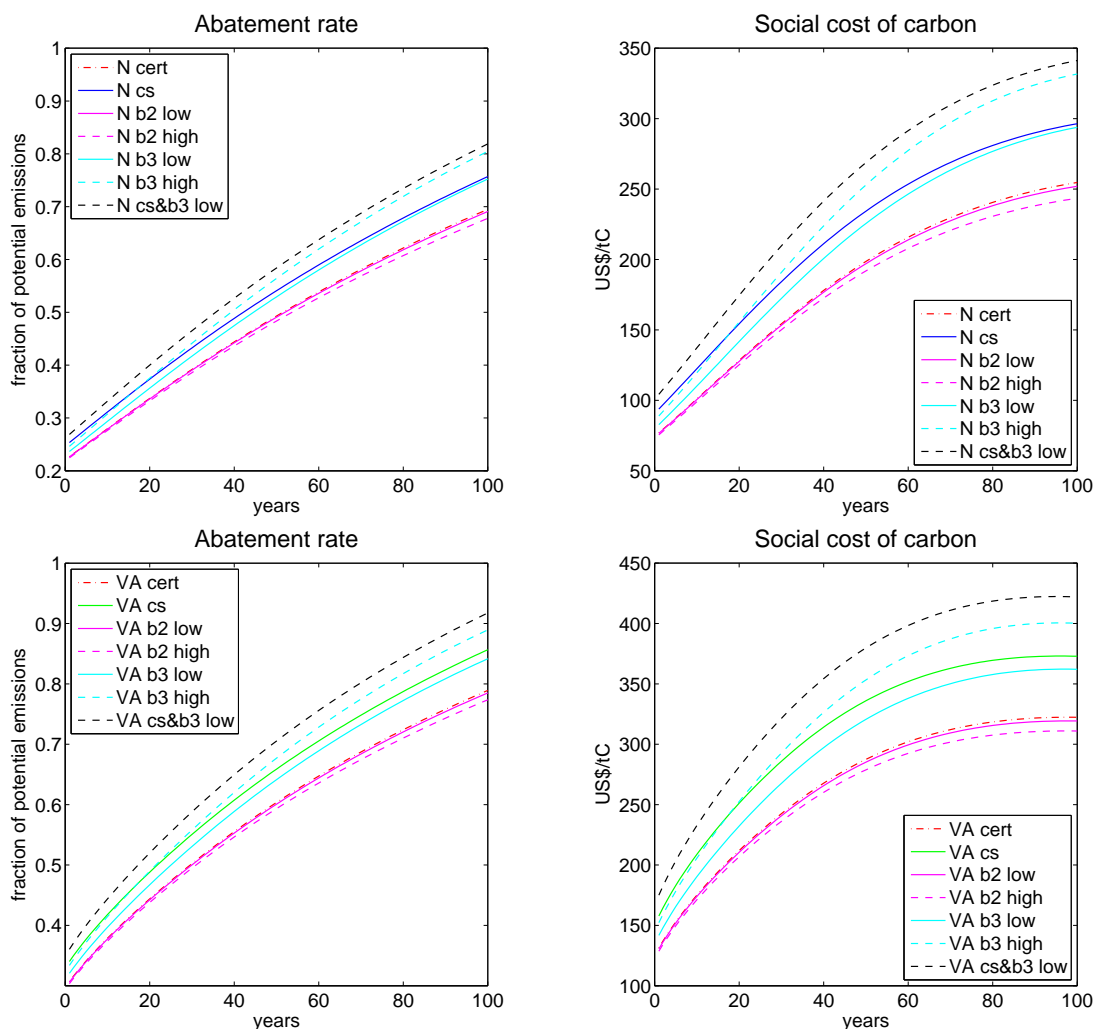


Figure 3 compares the optimal abatement rate and optimal carbon tax for different types of uncertainty for both, the scenario with $\eta = 2$ as in DICE-2007 (upper graphs) and for $\eta = \frac{2}{3}$ as suggested by Vissing-Jørgensen & Attanasio (2003) in the disentangled approach (lower graphs). “cs” denotes uncertainty over climate sensitivity, “b2” uncertainty over the damage coefficient, and “b3” uncertainty over the damage exponent. For details see section 3.2.

differences between the different scenarios.

If risk has an effect that is not due to risk aversion it has to be caused by other non-linearities in the underlying climate and economic system. Figure 3 compares three different types of risk, the lognormal distribution on climate sensitivity () and two different types of damage uncertainty. The first damage uncertainty is a normal distribution on the coefficient of the damage term (magenta) and the second places the normal on the exponent making damages quadratic only in expectation (cyan). The most striking result depicted in Figure 3 is that uncertainty on the damage coefficient b_2 reduces abatement as well as the social cost of carbon in both scenarios. Note that b_2 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 does not change much and can hardly be distinguished from the certain run in our model, the dotted version featuring the higher variance decreases abatement notably. On the contrary the normal distribution over the damage exponent (cyan) significantly increases abatement and the social cost of carbon for all scenarios. Recall that the variance of the b_3 uncertainty was chosen as to match the damages of the b_2 scenario approximately for one standard deviation at a 3°C temperature increase. The functional form how the three different uncertainties translate into a net production (and thus consumption and investment) loss is captured by equation (6) which can be rewritten as

$$Y_t^{net} = \frac{Y_t^*}{1 + b_2 T_t^{b_3}}, \quad (9)$$

where $Y_t^* = (1 - \Lambda(\mu_\tau)) Y_T^{gross}$, i.e the total production net the expenditure on abatement. b_2 uncertainty corresponds to a linear variation in the denominator of what becomes the net GDP. Thus, b_2 variation is translated by a convex function into variations of net GDP. Therefore, expected GDP loss under uncertainty over b_2 is actually lower than GDP loss using the expected coefficient. b_3 uncertainty also takes place in the denominator of equation (9), but in the exponent of the temperature. A straight forward calculation shows that the resulting transformation of b_3 into GDP loss is convex and, thus, corresponding uncertainty increases the expected GDP loss. Finally, uncertainty over climate sensitivity lies affects temperature. Temperature enters the denominator squared. However, that alone would not make the transformation translating uncertainty into GDP loss a convex function. However, climate sensitivity is a log-normal distribution. Thus, relating to normally distributed uncertainty it again takes place in the exponent. The remaining difference between uncertainty over climate sensitivity and over the damage exponent b_3 is observed by the following transformation of equation (9)

$$Y_t^{net} = \frac{Y_t^*}{1 + b_2 \exp[b_3(\ln s_t + \ln m_t)]},$$

with $m_t = \frac{\ln \frac{M_\tau}{M_{preind}}}{\ln 2}$ indicating the relevant measure of CO_2 stock increase. Comparing a normal distribution of b_3 to one of $\ln s_t$ we observe the only difference being that b_3 multiplicatively interacts with $\ln m_t$ in the exponent while climate sensitivity does not. Thus, we expect the effect of b_3 uncertainty to increase relatively to the effect of climate sensitivity uncertainty as the carbon stock increases. Indeed, we observe this pattern clearly in Figure 3 for all scenarios. We conclude from the results summarized in Figure 3 that the variance of the probability distribution capturing uncertainty clearly is relevant for determining optimal policies. However, we also conclude that the position at which we introduce uncertainty within the given functional estimated are even more relevant and might even change the direction of the effect that uncertainty has on the optimal carbon tax.

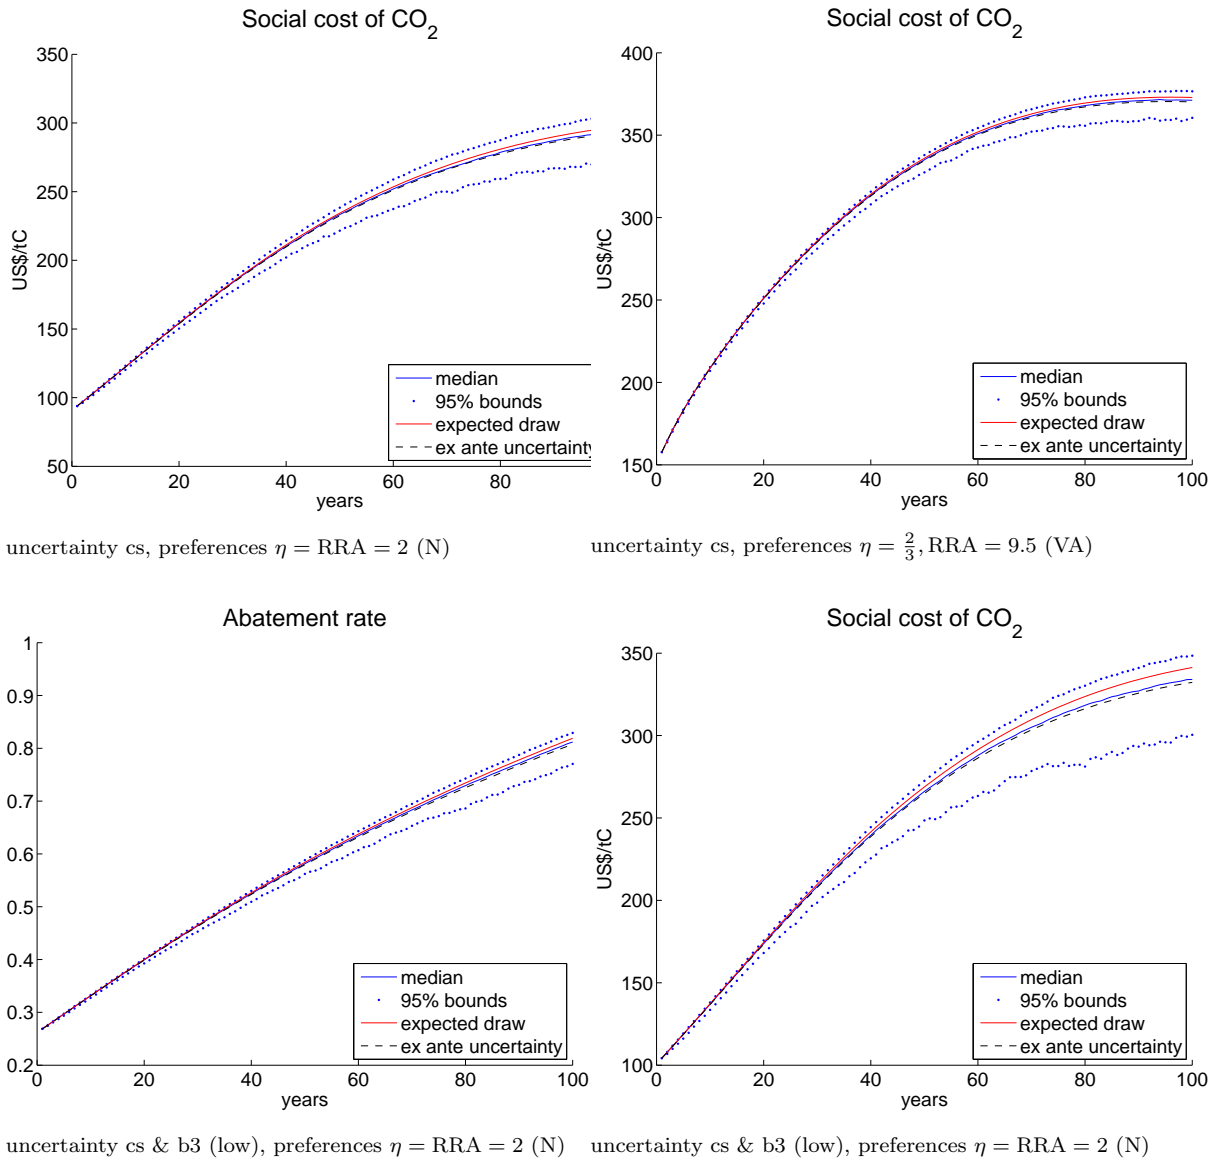


Figure 4 compares different paths characterizing the predicted time development under uncertainty.

So far our analysis compared optimal policy paths. As we pointed out in section 3.2 we generated these paths by using the optimal control rules under uncertainty and assuming that nature happens to draw expected values. In the following we discuss how our chosen path representation (expected draws) relates to other ways of constructing a representative time path under uncertainty. The most attractive alternative candidate is a Monte-Carlo simulation of individual paths. For this purpose we have randomly drawn the parameter realizations in every period for 1000 time paths for each scenarios. Figure 4 depicts the resulting median path as well as the boundaries of the interval containing 95% of the individual paths. We observe that the expected draws are slightly above the median. Moreover, the 95% bounds tell




Figure 5 compares the ex-ante uncertainty path as well as the median of 1000 randomly generated realizations with the path with obtain optimizing under uncertainty and nature drawing expected values. The left graph is based on $\eta = \text{RRA} = 2$ as in Nordhaus (2008) the right graph increases risk aversion to $\text{RRA} = 8.5$ (keeping $\eta = 2$).

us that, on average, downward deviations are larger than upward deviations. Figure 4 also depicts the optimal policy prediction with ex-ante uncertainty. This scenario does not employ the same value function as the earlier simulations. As explained in section 3.2 the ex-ante uncertainty scenario is based on drawing the uncertain parameters in period 0 and then using the control rules under certainty to generate the optimal policy paths. The depicted path represents the expected value over these individual paths. It corresponds to the method used in Nordhaus (2008) and most of its variations. The policy paths could become a realized “optimal” policy path under the constraint that policies have to be chosen once and for all in the first period. However, the expected real variables of an ex-ante uncertainty scenario generally cannot realize. Figure 5 shows an example that makes it particularly obvious that the expected ex-ante scenario is infeasible as a realization. The figure depicts the differences to the path we generated by assuming that nature draws the expected values, which is in particular a feasible path. For the first part of the century, the ex-ante uncertainty scenario goes along with a higher temperature and a lower damage than the path under expected draws. While quite attractive, such a path is unfortunately impossible as a realization. The graph is based on the “N” or $\eta = \text{RRA} = 2$ scenario. Despite the theoretical restrictions of the ex-ante method, Figure 4 shows that the prediction of the ex-ante uncertainty path is surprisingly close to both, the median and the path resulting from expected draws, both of which consider uncertainty striking in every period and optimization under uncertainty (as opposed to certainty in the ex-ante scenario).

Any individual path under uncertainty is of Lebesgue measure zero. Therefore, the most precise way to characterize and compare the different scenarios is by means of the corresponding control rules representing how the optimal policies react to

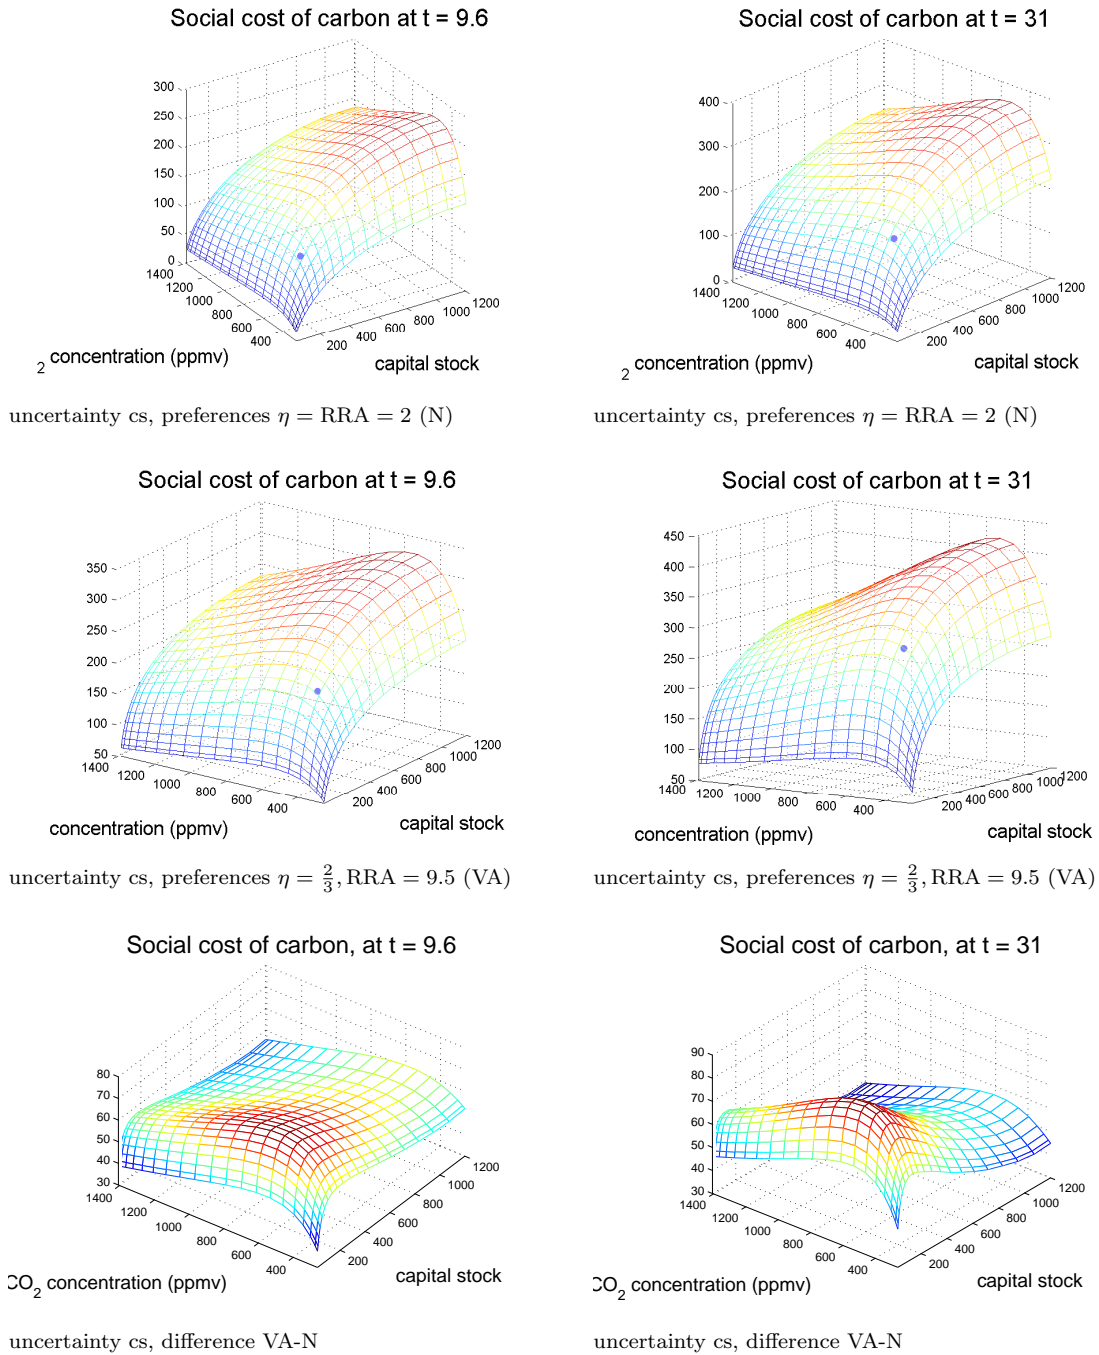


Figure 6 depicts and compares the optimal control rules for the scenarios based on Nordhaus (2008) preferences and those based on Vissing-Jørgensen & Attanasio (2003). We depict the control rule for $t = 9.6$ and for $t = 31$. The dots denote the optimal carbon tax at the respective point in time if nature realizes expected draws in all of the preceding periods (corresponding to the time paths presented in the previous figures).

uncertainty conditional on a particular stock variables (brought about by whatever

history of shock realizations). Figure 6 depicts the optimal control rules for the carbon tax in both, the Nordhaus and the VA scenarios. We evaluate the control rule at $t = 9.6$ years (time node 6) and $t = 50$ years (time node 10). The upper graphs represent the control rule for the setting based on Nordhaus's (2008) preferences, the graphs in the middle are based on the disentangled estimates of Vissing-Jørgensen & Attanasio's (2003), and the graphs at the bottom represent the differences between the two control rules in the respective settings. The dots mark the optimal carbon tax at the respective point in time if nature realizes expected draws in all of the preceding periods (corresponding to the time paths presented in the previous figures).¹⁰ We observe that in both scenarios the optimal carbon tax increases in the amount of capital in the economy, but at a decreasing rate. Note that the tax is measured relative to the marginal value of capital which decreases as capital increases. Mitigating a unit of carbon becomes than relatively more valuable as capital is more abundant. On the other hand, the carbon tax increases in CO_2 only at the beginning and (still around the level based on an optimal policy and expected realizations). However, for some concentration of CO_2 the optimal tax starts to fall again. Here the marginal (present value) damages of emitting another unit of CO_2 decreases again.¹¹ We also observe that the differences between the two scenarios in terms of the optimal carbon tax ranges between \$40-80 for a given state of the world (note that the path comparisons assumed different optimal policies and thus different states of the world in these periods). The highest differences between the two control rules are observed for in the region which corresponds closely to the points where either of the two control rules would bring us over the next decade or the next three decades. The differences keep to be as high for a higher carbon stock but decrease in the case of a faster growth of the capital stock.

5 Conclusions

In this paper we scrutinize the effects of risk in a DICE-like integrated assessment model. For this purpose we translate the DICE model of Nordhaus (2008) into a recursive dynamic programming framework, which allows us to analyze repeated stochasticity on a yearly basis. Most importantly, however, it allows us to disentangle risk preferences from the propensity to smooth consumption over time. We find that risk aversion has the least effect on the optimal abatement rate and the optimal carbon tax. In contrast, the propensity to smooth consumption over time has a major effect on both policies. A lower aversion to intertemporal substitution, as estimated in approaches that disentangle it from risk aversion, implies a significantly higher carbon tax than an evaluation based on the parameter specification used in Nordhaus (2008). A message is that for an appropriate evaluation in a DICE type

¹⁰Note that currently the time of the control rules is rounded to one digit while the "dots" represent the optimal tax at the rounded year. We will fix this imprecision in the next revision.

¹¹A result we plan to analyze further in a revised version of this paper.

framework we should pay more attention to get intertemporal substitutability right than to match risk aversion. In the standard model both coincide, and if we have to pick such a joint parameter, pick it to reflect aversion to intertemporal substitution.

While risk aversion is not a major determinant in our simulations, risk very well is. We have shown how damage uncertainty and uncertainty over the climate sensitivity can increase the optimal carbon tax and abatement rate significantly. However, it depends on how the risk is introduced into the system as the relevant non-linearities are found in the functional formulation of the interaction between CO₂ level and GDP loss. In particular, if we make the same coefficient in the damage function uncertain as Nordhaus (2008) we find that risk leads to a small decrease in the optimal carbon tax. However, if we make the other non-trivial parameter of the DICE damage function uncertain risk has a major effect on the optimal policies increasing the social cost of carbon significantly. While the current estimates of global climate damages are clearly surrounded by a significant amount of uncertainty, we cannot say which way of introducing uncertainty is the right way. The sensitivity of the optimal policies under risk to the particular functional forms emphasizes the necessity to further develop the functional forms estimating climate change damages taking into account how to capture the remaining uncertainties. For climate sensitivity we have better estimates surrounding the uncertainty and it is a priori more obvious how this uncertainty enters the integrated model. However, our analysis suggests¹² that a similar sensitivity of optimal policies as to damages holds with respect to the distributional assumption that we put on the climate sensitivity parameter. For our lognormal assumption informally based on a variety of estimates collected in IPCC (2001, Fig TS.25 p 65) we find significant effects of risk on both, abatement rate and social cost of carbon. From the related uncertainty scenarios we have analyzed, we hypothesize that a normal distribution as used in Nordhaus's (2008) ex-ante modeling would only have very minor effects. It will be the precise functional form of the distribution that matters more than the actual variance of capturing the uncertainty over climate sensitivity in an evaluation based on the DICE like integrated assessment.

We have criticized the ex-ante approach from a theoretical perspective and have shown that its predictions can generally lead to infeasible paths, featuring for example a higher temperature with lower damage than a realizable path for the setting based on Nordhaus's (2008) parameters. Despite this short-comings we have also shown that the prediction of an ex-ante model are effectively very close to that of our truly stochastic model. That finding gives a stronger stand to a variety of papers in the literature that have analyzed uncertainty in DICE like integrated assessment models based on ex-ante uncertainty. Ex-ante uncertainty can be interpreted as a setting with immediate resolution of uncertainty or perfectly correlated uncertainty. In that sense our setting is the complete opposite where uncertainty remains over the full time horizon. The truth lies inbetween, where uncertainty reduces over time as

¹²Here we infer from the sensitivity we have shown with respect to the damage coefficients and the analytical relation worked out in the results section. A simulation is to follow.

concentrations stabilize and as we learn over the underlying system. The similarity of the ex-ante uncertainty and our truly stochastic simulations suggest that a reduction of uncertainty over time per se is not likely to change the optimal policies much. However, extending our current risk sensitivity analysis to a setting including anticipated parametric learning will constitute a worthwhile extension complementing the current findings.

Appendix

A Notes on the Numerical Implementation

We use the collocation method described by Miranda and Fackler (2002) to numerically approximate the value function for our stylized climate model. This means that we approximate the value function by a linear combination of n Chebyshev polynomials:

$$\widehat{V}(K_t, M_t, S_t) = \sum_{j=1}^n c_j \phi_j(K_t, M_t, T_t) \quad (9)$$

here $\phi_1, \phi_2, \dots, \phi_n$ are the Chebyshev polynomials and c_1, c_2, \dots, c_n are parameters to be determined. The value function is approximated by finding the n coefficients, c_1, c_2, \dots, c_n , that satisfy the Bellman equation at n previously specified points in the state-space, called collocation nodes: $\text{beginEQA}[1] \sum_{j=1}^n c_j \phi_j(K_t, M_t, T_t) = \max_{x_t, \mu_t} \frac{c_t^\rho}{\rho} + \frac{\beta}{\rho} \left(\mathbb{E} \left[\rho \sum_{j=1}^n c_j \phi_j(K_t, M_t, T_t) \right]^\frac{\alpha}{\rho} \right)^\frac{\rho}{\alpha}$

This problem is solved using the following iterative algorithm: 1. Start with an initial guess $c_1^0, c_2^0, \dots, c_n^0$. 2. Given the current guess $c_1^k, c_2^k, \dots, c_n^k$, numerically solve the maximization problem $\text{beginEQA}[1] V_i^k = \max_{x_t, \mu_t} \frac{c_t^\rho}{\rho} + \frac{\beta}{\rho} \left(\mathbb{E} \left[\rho \sum_{j=1}^n c_j \phi_j(K_t, M_t, T_t) \right]^\frac{\alpha}{\rho} \right)^\frac{\rho}{\alpha}$ to get estimates, $V_1^k, V_2^k, \dots, V_n^k$, of the value function at the collocation nodes 3. Given the estimates $V_1^k, V_2^k, \dots, V_n^k$, find the parameters $c_1^{k+1}, c_2^{k+1}, \dots, c_n^{k+1}$ that solve the Bellman equation at the collocation nodes:

$$V_i^k = \sum_{j=1}^n c_j^{k+1} \phi_j(K_t, M_t, S_t), \forall i = 1, 2, \dots, n$$

4. Iterate steps 2. and 3. until the difference between $c_1^k, c_2^k, \dots, c_n^k$ and $c_1^{k+1}, c_2^{k+1}, \dots, c_n^{k+1}$ becomes small enough to satisfy a previously specified convergence criterion. The approximated value function can be used to estimate the optimal choices of consumption, investment and abatement as a function of current levels of capital, pollution and temperature. This decision rule can then be used to calculate the optimal time-path of the state and decision variables.




Figure 7 shows how we calibrated our exogenously decreasing decay rate of atmospheric carbon to DICE. The depicted model runs assume certainty and otherwise coinciding parameters as in the DICE-2007 model.

B Other Graphs

As we do not model the carbon cycle explicitly we assume that our rate of decay decreases exogenously over time, mimicking that reservoirs fill up over time. We calibrate our equation for the rate of decline of the decay of atmospheric carbon to the DICE-2007 model.¹³ Figure 7 depicts the carbon stock in the original DICE, in the one resulting from our calibration of the decay rate and for some alternative assumptions on decay. The left graph compares the results from DICE-2007 (decadal time step) and from our calibration to runs assuming a constant decay rate. The right graph compares the results from DICE-2007 and from our chosen calibration to alternative assumptions on the decline of the decay rates. δ_{M0} is the initial decay rate, $\delta_{M\infty}$ the asymptotic decay rate, and $g_{\delta M}$ denotes the rate at which the decay rate transition from its initial to its asymptotic value according to equation (4).

Figure 8 summarizes the exogenous drivers of the DICE model.

¹³We used the EXCEL version of the model that can be downloaded from William Nordhaus' website (<http://nordhaus.econ.yale.edu/DICE2007.htm>) as it generates a longer time series than depicted for example in Nordhaus (2008). Note that the EXCEL model assumes a constant savings rate. However, we found an almost constant savings rate in our model as well, and the match of EXCEL DICE to the full DICE seems to be a closer fit (for the first hundred years where we had the comparison) than we can get with our calibration.




Figure 8 summarizes the exogenous drivers of the DICE-2007 (and our) model.

References

- Ackerman, F., Stanton, E. A. & Bueno, R. (2010), 'Fat tails, exponents, extreme uncertainty: Simulating catastrophe in dice', *Ecological Economics* **69**(8), 1657–1665.
- 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.
- Giuliano, P. & Turnovsky, S. J. (2003), 'Intertemporal substitution, risk aversion, and economic performance in a stochastically growing open economy', *Journal of International Money and Finance* **22**(4), 529–556.
- 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.

- IPCC (2001), *Climate Change 2001: The Scientific Basis, Technical Summary*, Secretariat of the International Panel of Climate Change, Geneva. Edited by J.T. Houghton, Y. Ding, D.J. Griggs, M. Noguer and P.J. van der Linden and D. Xiaosu.
- Kelly, D. L. & Kolstad, C. D. (1999), 'Bayesian learning, growth, and pollution', *Journal of Economic Dynamics and Control* **23**, 491–518.
- 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.
- Rothschild, M. & Stiglitz, J. E. (1970), 'Increasing risk: I. A definition', *Journal of Economic Theory* **2**, 225–243.
- Traeger, C. (2007), 'Wouldn't it be nice to know whether Robinson is risk averse?'. Working Paper.
- 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.

Table 1 Parameters of the model

$\delta = 1.5\%$	discount rate
$\delta_K = 10\%$	depreciation rate of capital
$\delta_{M,0} = 1.7\%$	initial rate of decay of CO2 in atmosphere
$\delta_{M,\infty} = 0.25\%$	asymptotic rate of decay of CO2 in atmosphere
$\delta_M^* = 3\%$	initial rate of decay of CO2 in atmosphere
$\sigma = 0.1$	CO ₂ equivalent emission-GDP ration
$\kappa = .3$	capital elasticity in production
$a_1 = 0.12$	abatement cost, multiplicative constant
$a_2 = 2.8$	abatement cost, exponent
$b_1 = 0$	damage intercept
$b_2 = 0.00284$	coefficient of quadratic damage. For uncertain scenario normally distributed with standard deviation of 0.0025.
$b_3 = 2$	damage exponent. For uncertain scenario normally distributed with standard deviation of 0.35.
$s = 3.08$	Climate sensitivity, i.e. equilibrium response to doubling of atmospheric CO ₂ concentration with respect to preindustrial concentrations. For uncertain scenario lognormally distributed with $\mu = 1$ and $\sigma = .5$ implying an expected value of 3.08 and a standard deviation of approximately 2.7.
$M_{preind} = 596$	in GtC, preindustrial stock of CO2 in the atmosphere
$K_0 = 137$	in trillion 2005-USD, initial value for global capital stock
$M_0 = 80.9$	in Gt, Initial stock of atmospheric CO ₂
$L_0 = 6514$	in millions, population in 2005
$L_\infty = 8600$	in millions, asymptotic population
$\delta_L^* = 0.35$	rate of convergence to asymptotic population
$g_{\sigma,0} = -0.73\%$	initial rate of decarbonization
$\delta_\sigma = 0.3\%$	initial rate of decarbonization
$g_{A,0} = 0.2722\%$	initial rate of decarbonization
$\delta_A = 1\%$	initial rate of decarbonization
$B_0 = 11$	in Gt, initial CO2 emissions from LUCF
$g_B = -1\%$	growth rate of CO2 emisison from LUCF
$a_0 = 1.17$: 1.17, cost of backstop 2005
$a_1 = 2$	ratio of initial over final backstop cost
$a_2 = 2.8$	cost exponent
$g_\Psi = -0.5\%$	rate of convergence from initial to final backstop cost