Climate Policy and Growth Uncertainty: Dicing with DICE[1]

Svenn Jensen[†] & Christian Traeger[§]

[†] Ragnar Frisch Centre for Economic Research, Gaustadalleen 21, 0349 Oslo, Norway; Statistics Norway, Kongens gate 6, 0153 Oslo, Norway; svenn.jensen@frisch.uio.no

[§] Department of Agricultural & Resource Economics, UC Berkeley 207 Giannini Hall #3310, Berkeley, CA 94720-3310, USA; traeger@berkeley.edu

Version of August 2012

Abstract: Integrated assessments of climate change commonly rely on the assumption that technological progress outgrows climate change damages by an order of magnitude, even without any climate policy. Then, mitigating greenhouse gases is a redistribution from the poor present to a rich future. The optimal climate policy is highly sensitive to these growth assumptions. While the world economy experienced enormous growth over the last century, it remains uncertain whether such growth can be sustained over several more centuries. We incorporate growth uncertainty into an integrated assessment model that was recently employed to determine the US federal social cost of carbon. We derive optimal carbon taxes and mitigation rates in a stochastic dynamic programming framework, solving the non-linear, out-of-steady state problem. This approach differs largely from the Monte-Carlo simulations that are current state of the art in the integrated assessment of climate change. The standard intertemporally additive expected utility model falls short of simultaneously capturing risk premia and risk-free discount rates (equity premium puzzle, risk-free rate puzzle). The finance literature shows that fully rational Epstein-Zin-Weil preferences, which disentangle risk aversion from the propensity to smooth consumption over time, resolves these shortcomings. We derive optimal climate policy under standard preferences and under comprehensive Epstein-Zin-Weil preferences analyzing how these policies respond to long-term growth risk. Our findings suggest that growth uncertainty can change today's optimal social cost of carbon by up to 23 %, and that risk aversion is a major determinant of optimal climate policy. Moreover, the sign of the risk effect switches with changes in the propensity to smooth consumption. We explain this effect analytically in terms of prudence and pessimism effects.

JEL Codes: Q54, Q00, D90, D8, C63

Keywords: climate change; integrated assessment; social cost of carbon; uncertainty; growth; risk aversion; intertemporal risk aversion; precautionary savings; prudence; Epstein-Zin preferences; recursive utility; dynamic programming; DICE

[1]We thank Benjamin Crost, Reyer Gerlagh, Michael Hoel, Bard Harstad, Alfonso Irarrazabal, Larry Karp, Derek Lemoine and Tony Smith for help and insight.

1 Introduction

Future economic growth is of first order importance for climate change evaluation. Extrapolating growth from the past century to the coming centuries makes greenhouse gas mitigation a redistribution from the present poor to the future rich. For example, even in the absence of any climate change policy, Nordhaus's (2008) widespread DICE-2007 model implies that generations living 100 years from now are five times richer than today's generation. We analyze how uncertainty about economic growth affects optimal climate policy. We model fundamental uncertainty about technological progress that is independent of climatic change. This growth uncertainty can also be interpreted as a consequence of economic crises, speading of stable social institutions, or social unrest and war. We do not model a direct impact of climate change on economic growth. While such a direct link would have a major impact on economic policy, this direct link is empirically more controversial than the fundamental growth uncertainty we depict. Our paper is the first to consistently analyze how growth uncertainty impacts optimal climate policies in the integrated assessment of climate change. We focus on optimal abatement effort and the optimal carbon tax, but we also discuss how capital investment reacts optimally to the uncertainty. We employ a recursive dynamic programming version of the DICE-2007 model by Nordhaus (2008). This model is the most widespread integrated assessment model and was recently used as one of three models for determining the US federal social cost of carbon (Interagency Working Group on Social Cost of Carbon 2010).

It is widely known that the standard economic model is not able to simultaneously capture observed risk premia and discount rates. Agents have a significantly higher willingness to pay for risk avoidance than common parameterizations of this intertemporally additive expected utility model suggest (equity premium puzzle). Increasing risk aversion, however, would simultaneously increases aversion to intertemporal substitution implying excessive discount rates (risk-free rate puzzle). An important branch of the finance literature resolves this puzzle by introducing Epstein-Zin-Weil preferences in combination with persistent shocks (Vissing-Jørgensen & Attanasio 2003, Bansal & Yaron 2004, Bansal et al. 2010, Nakamura et al. 2010, Chen et al. 2011). Epstein-Zin-Weil preferences disentangle risk attitude from the propensity to smooth consumption over time. There is no a priori reason why these different preference characteristics should coincide, and Epstein-Zin-Weil preferences satisfy the usual rationality assumptions including time consistency and the von Neumman-Morgenstern axioms (Traeger 2010). We analyze the implication of uncertainty under standard preferences as well as under these comprehensive Epstein-Zin-Weil preferences that improve the capital market calibration of the DICE-2007 model.

Two recent papers address technological uncertainty in a dynamic climate change context. Fischer & Springborn (2011) investigate the optimal labor and output responses to short term productivity shocks (business cycles) under three climate policy instruments: a tax, a cap and trade program, and an intensity target. Their model features two state variables, productivity and capital, and they abstract from the climate system and assume an exogenous emission target each period. They argue that no policy is strictly preferred from the others. Heutel (2011) considers how optimal climate policy reacts to short term productivity shocks. His model features technology, capital and carbon as stocks. The shock is modeled by a mean reverting Markov process. He calibrates his model to analyze unilateral climate action by the US, therefore emissions outside the US are treated as exogenous and constant. He finds that optimal emissions are procyclical, but dampened relative to a laissez-faire economy. Both papers consider short term fluctuations, whereas we are interested in the uncertainty about the growth trend. Moreover, the authors assume stationarity and use (log-)linearizations around the steady state to solve their models. Neither of the models are hence typical climate-economy models, which are mainly concerned with off-equilibrium dynamics. Stochastic growth processes also characterize dynamic stochastic general equilibrium (DSGE) models in macroeconomics. Whereas the original real business cycle models focused on fluctuations around a trend, i.e. mean reverting processes, the more recent literature also considers fluctuations in the trend, or random walks. For example, in Aguiar & Gopinath (2007) variations in growth trend stochasticity explain the differences in business cycle characteristics in developing and developed countries. The central methodological difference compared to our model is again the use of linearizations around the steady state as a solution method. Baker & Shittu (2008) review the climate change literature on endogenous technological change and uncertainty. This strand of the literature models uncertainty in climate technology, not uncertainty about the general productivity in the economy.

A different strand of literature applies Epstein-Zin-Weil preferences to the analysis of climate policy. The only paper concerned with growth uncertainty is a highly stylized analytic model of the social discount rate by Traeger (2012). He explains the importance of a comprehensive risk evaluation and shows that uncertainty has a negligible impact on discounting in the standard model. Epstein-Zin-Weil preferences reduce the discount rate for two accounts. First, the disentangled estimate of the intertemporal elasticity of substitution is higher, which reduces the risk-free rate. Second, disentangled risk aversion implies a major reduction of the discount rate, where the correlation between policy payoffs and baseline growth become highly relevant. The paper neither models the climate system nor does it distinguish between produced capital and environmental capital. Crost & Traeger (2010) use Epstein-Zin-Weil preferences in a recursive version of DICE to evaluate damages. They show that the disentanglement is of major importance for long-term evaluation because it gets the risk-free discount rate right. They find that risk aversion itself has virtually no effect on the optimal policy under damage uncertainty. In contrast, we find that risk aversion has a significant impact on optimal policies in the context of growth uncertainty. In a more stylized mitigation model without explicit temperature dynamics and with only two states of the world, Ha-Duong & Treich (2004) analyze damage uncertainty and find that risk aversion and aversion to intertemporal substitution can affect policy in opposite directions.

Employing standard preferences, Kelly & Kolstad (1999) and Leach (2007) analyze learning about the sensitivity of global temperatures to CO_2 in a recursive implementation of DICE. Lemoine & Traeger (2010) and Lontzek et al. (2012) analyze the policy impact of tipping points in the climate system in a similar DICE

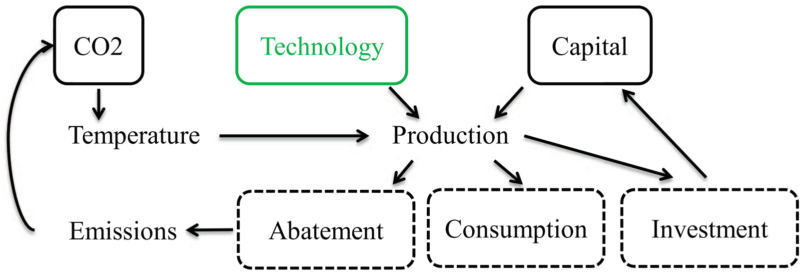

Figure 1 is an abstract representation of the climate-enriched economy model. The control variables consumption and abatement as well as the 'residual' investment are represented by dashed rectangles. The main state variables are depicted by solid rectangles. The green color indicates that the technology level is uncertain.

implementation. Keller et al. (2004) model uncertainty and learning with a simplified uncertainty tree in a non-recursive implementation of DICE. Webster et al. (2012) analyze the impact of uncertainty in the cost of abatement in a stochastic model based on DICE-99, finding it has a negligible effect. In their model, contrary to ours, capital investments are exogenous and the time horizon is finite. They reduce the resolution of their model using 50 year time intervals, and solve it by approximate dynamic programming techniques. Cian & Tavoni (2011) model technological uncertainty in a non-recursive two stage uncertainty tree implementation of the IAM WITCH. More remotely related are a set of Monte-Carlo simulations that approximate uncertainty by averaging over deterministic runs of IAMs (Hope 2006, Nordhaus 2008, Ackerman et al. 2010, Anthoff & Tol 2010, Richels et al. 2004, Dietz 2009).

2 Model and welfare specification

Integrated assessment models embed a model of the world economy in a model of the climate system to investigate their interactions. We build a recursive version of the DICE-2007 model, with some minor simplifications.[2] Our model is summarized graphically in Figure 1. The world economy is governed by the classic Ramsey-Cass-Koopmans growth model. Technology level and population evolve exogenously, while capital accumulation is endogenous. Production of an aggregate commodity causes emissions that accumulate in the atmosphere. The social planner can spend part of the production on emission reductions (abatement). The emission stock in the atmosphere changes the Earth's energy balance and causes global warming. An increase of global average temperature above the level of 1900 causes damages that reduce world output. We solve for the optimal investment and abatement decisions.

[2]In order to avoid the "curse of dimensionality" in our infinite horizon dynamic programming version of DICE, we replace the carbon cycle in DICE by single decay rate fit, and we simplify the equation of motions for temperatures (Appendix D). The simplified model is calibrated to perfectly fit the baseline policies in DICE, but temperatures are slightly lower than in the original model.

2.1 Growth Uncertainty

The rate of technological progress is uncertain. The technology level enters the Cobb-Douglas production function and determines the overall productivity of the economy. A shock in the growth rate permanently affects the technology level in the economy. The technology level A_t in the economy follows the equation of motion[3]

$$\tilde{A}_{t+1} = A_t \exp[\tilde{g}_{A,t}] \quad \text{with} \quad \tilde{g}_{A,t} = g_{A,0} * \exp[-\delta_A t] + \tilde{z}_t \ . \tag{1}$$

The deterministic part of the stochastic growth rate $\tilde{g}_{A,t}$ decreases over time at rate δ_A as in the original DICE-2007 model.[4] We add a stochastic shock $\tilde{z}$, which is either identically and independently distributed (iid) or persistent.

Our main set of simulations analyzes the consequences of an iid shock

$$\tilde{z}_t \sim \mathcal{N}(\mu_z, \sigma_z^2) \ .$$

We base this value somewhat loosely on Kocherlakota's (1996) observation for the last century of US data that the standard deviation of consumption growth is about twice its expected value and set the standard deviation at twice the initial growth rate ($\sigma_z = 2 * g_{A,0} \approx 2.6\%$).[5] We fix the mean of the growth shock so that future expectations for the technology level coincide with those under certainty.[6] Figure 2 illustrates the future technology level under expected growth in solid green, and the 95% (simulated) confidence interval under iid growth shocks in dashed blue. In expectation, and in the deterministic model, the productivity level of the economy increases roughly threefold over the 100 year time horizon.

In a modification, we analyze the consequences of persistence in the growth shock. While our shocks always have a persistent effect on the technology level, persistence in the growth shock implies that the growth rate itself is intertemporally correlated. Persistent shocks are employed by the finance literature explaining the equity premium and the risk-free rate puzzle (Bansal & Yaron 2004). Here, we think of the persistent shock as a more fundamental uncertain change affecting technological progress, e.g. times of economic crisis, international conflict, or simply fundamental innovations or their absence. The theoretical literature has established that persistent shocks imply decreasing social discount rates over time (Weitzman 1998, Azfar 1999, Newell &

[3]Our numerical values correspond to the more widely used labor-augmenting formulation of technological progress. Given Cobb-Douglas production, it is formally equivalent to Nordhaus (2008) formulation, but leads to balanced growth also in the case of more general production specifications.

[4]We approximate all exogenous processes in DICE by their continuous time dynamics and evaluate them at a yearly step.

[5]The rate of technological progress drives consumption growth in the Ramsey-Cass-Koopmans economy. Our decision maker can smooth the effect of technology shocks using capital to smooth consumption. Moreover the steady state consumption growth rate also depends on deterministic population growth. Thus, our model is not build to reproduce or calibrate consumption fluctuations. We merely take the above reasoning as a proxy for a relevant order of magnitude.

[6]A mean zero shock of the growth rate would, by Jensen's inequality, imply an increase in the expected next period technology level. The technology level in period $t+1$ is determined by the random variable $\exp[\tilde{z}]$ that is lognormally distributed. Setting $\mathbf{E}[\tilde{z}] = -\sigma^2(\tilde{z})/2$ implies $\mathbf{E}\exp[\tilde{z}] = 1$ and that the expected technology level equals its deterministic part. A short calculation shows that the by the same reasoning $A_{t+\tau}$ expectations coincide with the deterministic part for all $\tau > 0$.

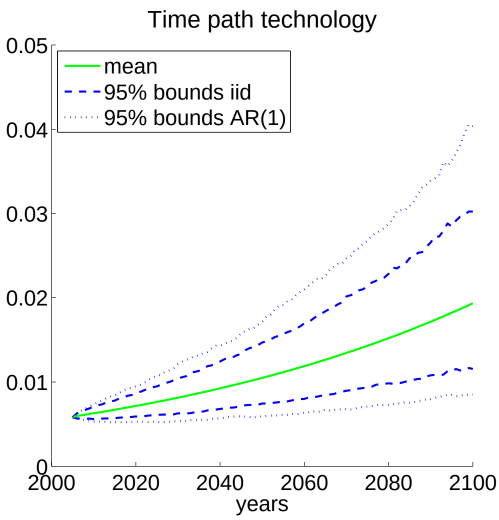

Figure 2 shows the expected draw and the 95% confidence intervals for technology time paths based on 1000 random draws of technology shock $\tilde{z}$ time paths with $\sigma_{\tilde{z}} = 2 * g_{A,0}$. The dotted lines are the confidcence intervals for an AR(1), while the dashed lines correspond to iid shocks.

Pizer 2003). We model persistence in form of an AR(1) process

$$\begin{aligned} \tilde{z}_t &= \tilde{x}_t + \tilde{y}_t \quad \text{where} \\ \tilde{x}_t &\sim \mathcal{N}(\mu_x, \sigma_x^2) \quad \text{and} \\ \tilde{y}_t &= \zeta y_{t-1} + \tilde{\epsilon}_t \text{ with } \tilde{\epsilon}_t \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2) \ . \end{aligned} \tag{2}$$

Choosing the standard deviations $\sigma_x = \sigma_\epsilon = \sqrt{2} * g_{A,0}$ once more results in a standard deviation of the overall shock $\tilde{z}_t$ of twice the initial growth rate. Our second specification coincides with the first in the case of vanishing persistence $\zeta = 0$, and positive persistence increases long-run uncertainty. We fix the mean values by requiring that the expected technology path once more corresponds with the one under certainty, conditional on $y_t = 0$.[7] Our simulations assume that 50% of the ϵ-shock carries over to the growth rate of the next year: $\zeta = 0.5$. The dotted lines in Figure 2 represent the 95% (simulated) confidence intervall for the technology levels over the next 100 years under such persistent growth shocks. While modeling an even higher persistence would be desirable, a random walk in the growth rate (instead of a mean reverting process) is a serious numerical challenge in an infinite horizon dynamic programming model. Persistence of the shock adds significantly to this challenge. We will show that even our rather moderate persistence has clear implications for optimal climate policy.

[7] A short calculation shows that we achieve this equivalence by setting $\mathbf{E}[\tilde{x}] = \mathbf{E}[\tilde{\epsilon}] = -\sigma^2(\tilde{x})/2$. An iteration of this reasoning shows that indeed $\mathbf{E}(A_{t+\tau}|y_t = 0) = A_{t+\tau}^{det}$, where $A_{t+\tau}^{det}$ is the technology level τ periods in the future when initializing the deterministic model in period t with the current technology level.

2.2 Welfare and Bellman equation

The decision maker maximizes her value function subject to the constraints imposed by the climate-enriched economy. We formulate the decision problem recursively using the Bellman equation. This recursive structure facilitates the proper treatment of uncertainty and the incorporation of comprehensive risk preferences. The relevant physical state variables describing the system are capital K_t, atmospheric carbon M_t, and the technology level A_t. In addition, time t is a state variable that captures exogenous processes including population growth, changes in abatement costs, non-industrial GHG emissions, and temperature feedback processes. Finally, in the case of persistent shocks, the state d_t captures the persistent part of last period's shock that carries over to the current period. We first state the Bellman equation for standard preferences, i.e. the time additive expected utility model:

$$V(K_t, M_t, A_t, t, d_t) = \max_{C_t,\mu_t} \frac{L_t\left(\frac{C_t}{L_t}\right)^{1-\hat{\eta}}}{1-\hat{\eta}} + \exp[-\delta_u]\mathbf{E}\left[V(K_{t+1}, M_{t+1}, \tilde{A}_{t+1}, t+1, \tilde{d}_{t+1})\right]. \tag{3}$$

The value function V represents the maximal welfare that can be obtained given the current state of the system. Utility within a period corresponds to the first term on the right hand side of the dynamic programming equation (3). It is a population (L_t) weighted power function of global per capita consumption (C_t/L_t). The parameter $\hat{\eta}$ captures two preference characteristics: the desire to smooth consumption over time and Arrow-Pratt relative risk aversion. Following Nordhaus (2008), we set $\hat{\eta} = 2$. The second term on the right hand side of equation (3) represents the maximally achievable welfare from period $t+1$ on, given the new states of the system in period $t+1$, which follow from the equations of motion summarized in Appendix D. The planner discounts next period welfare at the rate of pure time preference $\delta_u = 1.5\%$ ("utility discount rate"), again taken from Nordhaus's (2008) DICE-2007 model. In period t, uncertainty governs the realization of next period's technology level $\tilde{A}_{t+1}$ and, thus, gross production. Therefore, the decision maker takes expectations when she choses the optimal control variables consumption C_t and abatement rate μ_t (in DICE: emission control rate). The abatement rate characterizes the fraction of business as usual emissions that are mitigated because of climate policy. Equation (3) states that the value of an optimal consumption path starting in period t has to be the maximized sum of the instantaneous utility gained in that period and the welfare gained from the expected continuation path. The control C_t balances immediate consumption gratification with the value of future capital. The control μ_t balances immediate consumption that is given up for abatement against the reductions of future atmospheric carbon.

The standard model underlying equation (3) assumes that intertemporal choice over time also determines risk aversion, and the single parameter $\hat{\eta}$ governs both relative risk aversion and aversion to intertemporal change. However, a priori these two preference characteristics are distinct and forcing them to coincide implies the well-known equity premium and risk-free rate puzzles. Translated to climate change evaluation, these puzzles tell us that a calibration of standard preferences to asset markets, as done for DICE-2007, will result in a model that overestimates the discount rate and underestimates risk aversion. Epstein & Zin (1989) and Weil (1990) show how to disentangle the two, and Bansal & Yaron (2004) show how this disentangled approach resolves the risk-free rate and the equity premium puzzles. We emphasize that the model satisfies the usual rationality constraints including time consistency and the von Neumann & Morgenstern (1944) axioms, and it is normatively no less desirable than the standard discounted expected utility model (Traeger 2010). The latter paper also shows how to shift the non-linearity from the time-step as in Epstein & Zin (1989) to uncertainty aggregation, resulting in the Bellman equation

$$V(K_t, M_t, A_t, t, d_t) = \max_{C_t,\mu_t} \frac{L_t\left(\frac{C_t}{L_t}\right)^{1-\eta}}{1-\eta} \tag{4}$$
$$+ \quad \frac{\exp[-\delta_u]}{1-\eta} \left(\mathbf{E}\left[(1-\eta)V(K_{t+1}, M_{t+1}, \tilde{A}_{t+1}, t+1, \tilde{d}_{t+1})\right]^{\frac{1-\mathrm{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\mathrm{RRA}}} .$$

The parameter η captures the desire to smooth consumption over time. It measures aversion to intertemporal substitution and is, thus, the inverse of the intertemporal elasticity of substitution. The parameter RRA depicts the Arrow-Pratt measure of relative risk aversion. In the case $\eta = \mathrm{RRA}$ equation (4) collapses to equation (3). For a detailed analysis of the interpretation of the parameters RRA and ρ we refer to Epstein & Zin (1989) and to Traeger (2010). We base our choices of values for the disentangled preference on estimates by Vissing-Jørgensen & Attanasio (2003), Bansal & Yaron (2004), and Bansal et al. (2010). These papers suggest best guesses of $\eta = \frac{2}{3}$ and of relative risk aversion in the proximity of the value $\mathrm{RRA} = 10$ that we adopt. The social cost of carbon in current value units of the consumption-capital good is the ratio of the marginal value of a ton of carbon and the marginal value of a unit of the consumption good: $SCC_t = \frac{\partial_{M_t} V}{\partial_{K_t} V}$.In our optimization framework, the social cost of carbon is the optimal carbon tax.

2.3 Numerical Implementation

We give a brief summary of the numeric implementation, discussing details in Appendix B. We approximate the value function by Chebychev polynomials and solve the Bellman equation by value function iteration. We represent the continuous distribution capturing technological progress by Gauss-Legendre quadrature nodes. The Bellman equations (3) and (4) are not convenient for a numerical implementation for two reasons. First, capital and technology are subject to enormous growth and any value function approximation with a reasonable number of nodes would be rather coarse on the state space.[8] Second, modeling a random walk without mean reversion

[8]More precisely, the relevant part of the state space at different times would be disconnected. Our renormalization achieves that the relevant values lie in the same reduced region of the state space at all times. That allows us to obtain a much better approximation of the value function with less nodes.

is a major challenge and the Bellman equation as cited above would not converge with the amount of uncertainty we are capturing. Therefore, we renormalize consumption and capital in per effective labor units. For the technology level, our state variable captures the deviation from the deterministic evolution of technology. Finally, we map the infinite time horizon on a $[0, 1]$ interval. We take advantage of these change conveniently reformulating the Bellman equation to the form stated in equation (8) in Appendix B.

3 Results

We first illustrate the small impact of uncertainty in the entangled standard model. We then increase the disentangled coefficient of relative risk aversion to the value suggested in the finance literature and, in a second step, introduce persistence to the growth shock. Finally, we analyze the dependence of optimal policy on the propensity to smooth consumption over time.

3.1 Entangled standard preferences ($\eta = \text{RRA} = 2$)

Figure 3 presents optimal policies in the standard model ($RRA = \eta = 2$). The green lines present the optimal policies if the decision maker employs a deterministic model with expected growth rates. The dashed blue lines present the optimal policies in the presence of the iid growth shock discussed in section 2.1 that makes the technology level a random walk. Here, the decision maker optimizes under uncertainty, but nature happens to still draw expected values.[9] Stochasticity of economic growth implies a very minor increase in optimal mitigation and the corresponding carbon tax. For the current century, the optimal abatement is .2-.6 percentage points higher under uncertain than under certain growth. The optimal carbon tax increases between \$1 and \$4.5. In addition, current investment goes up by .35 percentage points. Hence, we find a small precautionary savings effect in both capital dimensions: produced productive capital and natural climate capital. In his analysis of the social discount rate, Traeger (2012) explains the smallness of the precautionary effect by pointing out that decision makers with entangled preferences are effectively "intertemporal risk neutral".[10]

[9] The optimal policy in period t depends on growth realizations up to period t. Our actual solution derives control rules that depend on all states of the system. Our path representation in Figure 3 makes actual growth identical to the deterministic case and fleshes out the policy difference arriving only from acknowledging uncertainty when looking ahead. We compare this representation to other possibly path represenations in Figure 8 in Appendix A.

[10] See Traeger (2012) for a detailed discussion and Traeger (2010) for the axiomatic foundation of intertemporal risk aversion. The basic idea of intertemporal risk neutrality is the following. Suppose a decision maker is indifferent between two consumption paths both of which fluctuate over time. From these two paths, construct a "high" consumption path by picking the higher consumption outcome in each period, and a "low" consumption path by picking the lower consumption outcome in each period. The decision maker is intertemporal risk neutral if he is indifferent between receiving either of the two original paths with certainty and receiving a lottery with a 50/50 chance over the "high" and the "low" path.

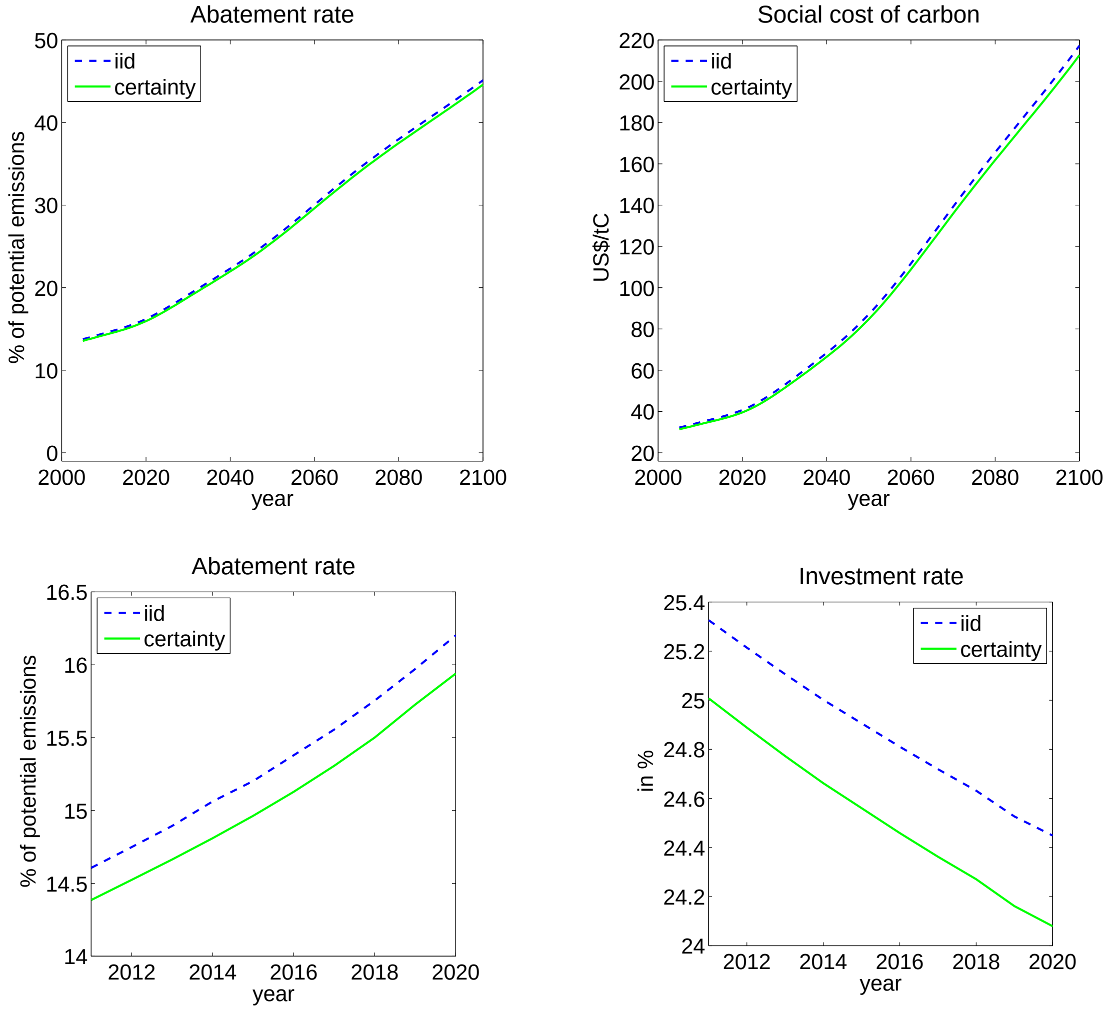

Figure 3 compares the optimal abatement rate, the social cost of carbon and the investment rate under certainty and iid uncertainty with standard preferences and RRA $= \eta = 2$.

3.2 Increasing risk aversion (RRA $= 10$) and persistence

The standard model of the previous section does not accurately capture risk premia (equity premium puzzle). The finance literature explains observed risk premia by increasing the (disentangled) coefficient of relative risk aversion to values around RRA $= 10$ (Vissing-Jørgensen & Attanasio 2003, Bansal & Yaron 2004, Bansal et al. 2010, Nakamura et al. 2010). Similarly increasing aversion to intertemporal substitution would imply much too high discount rates (risk free rate puzzle, see also graphic illustration in Traeger (2012) in the social discounting context). We therefore improve the DICE-2007 calibration to asset markets by employing Epstein-Zin-Weil preferences in the disentangled Bellman equation (4) with higher risk aversion.

Figure 4 shows the optimal climate policy keeping aversion to intertemporal substitution fix at $\eta = 2$ and increasing Arrow-Pratt risk aversion to $RRA = 10$. We observe a modest increase in abatement under uncertainty. The optimal abatement rate in 2012 increases by 12% to 16 percentage points. The optimal present day carbon tax increases by 23% to \$43 per ton of carbon. Similarly, the investment in produc-

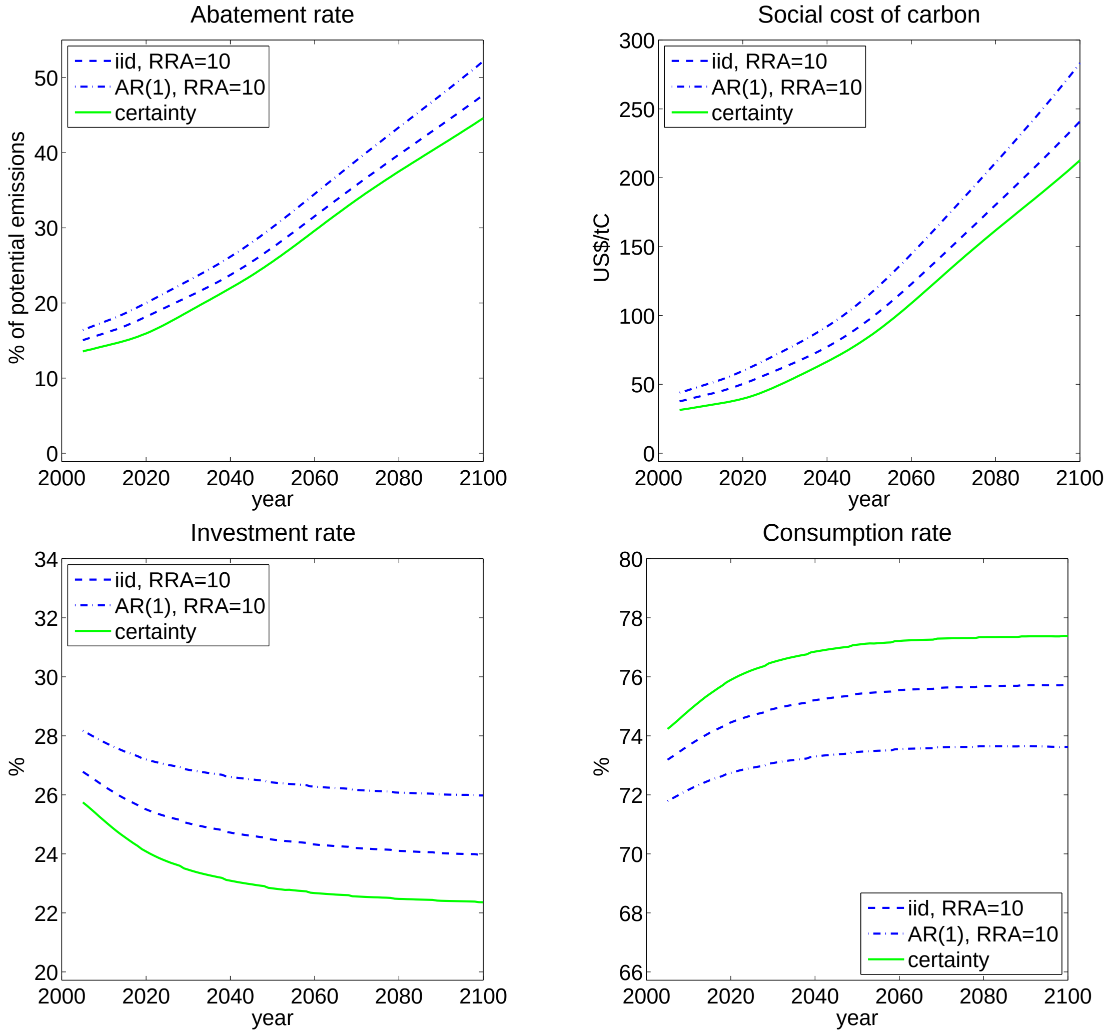

Figure 4 compares the optimal abatement rate, social cost of carbon, investment and consumption under certainty, an iid shock and a persistent shock with Epstein-Zin preferences, a coefficient of relative risk aversion of RRA = 10 and a coefficient of aversion to intertemporal substitution of $\eta = 2$.

tive capital increases. The more risk averse decision maker is more cautious, abating and investing more and consuming less. Robustness checks (not shown) confirm that these effects increase in the variance of the stochastic shock. With Arrow-Pratt risk aversion exceeding the consumption smoothing parameter ($RRA = 10 > \eta = 2$), the decision maker is now intertemporal risk averse.

The iid growth shocks have a permanent impact on the technology level, making technology a random walk. These iid shocks, however, do not capture that technological progress is intertemporally correlated. We therefore model a relatively moderate persistence of growth shocks according to equation (2). In addition to an iid shock component, the rate of technological growth experiences a persistent shock whose impact on technological growth decays by 50% per year.

The dashed-dotted lines in Figure 4 show the optimal climate policy under persistent growth shocks. Introducing persistence amplifies the long-run uncertainty, while keeping immediate uncertainty unchanged. Our moderate persistence in the shock approximately doubles the impact of uncertainty on optimal climate policy. The optimal abatement rate in 2012 increases by 24% to 18 percentage points, and the optimal carbon tax increases by 45% to $51 (both percentage increases with respect to the deterministic case).

3.3 Decreasing consumption smoothing

A further step in improving the DICE-2007 calibration to observed interest rates and asset returns is a more careful analysis of agents' propensity to smooth consumption over time. The value of $\eta = 2$ chosen in DICE implies an excessive risk-free discount rate and the finance literature shows that a reduction of the (disentangled) consumption smoothing parameter to values around $\eta = 2/3$ explains observed asset prices significantly better (Vissing-Jørgensen & Attanasio 2003, Bansal & Yaron 2004, Bansal et al. 2010, Nakamura et al. 2010, Chen et al. 2011).[11] The solid lines in Figure 5 display the effect of lowering η from 2 to 2/3 under certainty. The reduction in the parameter and, thus, the risk-free discount rate increases optimal mitigation significantly. The optimal carbon tax more than doubles (from $35 to $85 in 2012) and the optimal abatement rate close to doubles (from 14.5 to 24 percentage points in 2012). The decision maker is now less averse to shifting consumption over time. Hence, she evaluates the prospect of additional welfare for the relatively affluent generations in the future more positively than a decision maker with a higher propensity to smooth consumption. Crost & Traeger (2010) also point out this effect, which does not depend on the uncertain growth.

The dashed lines Figure 5 represent optimal policy under growth uncertainty, when $\eta = 2/3$, $RRA = 10$, and the shock is iid. The optimal abatement and the social cost of carbon fall over the full time horizon. The sign of the uncertainty effect is opposite to the one observed in the earlier settings. Its magnitude, however, is of similar to the case with $\eta = 2$: abatement in 2012 decreases by 9% to 22 percentage points. In contrast, investment in man-made capital still increases. The investment rate goes up by 2% (as opposed to 5% for $\eta = 2$), implying an optimal investment rate of almost 31% in the present but declining over time. Similarly, the consumption rate continues to decrease under uncertainty. Observe that the abatement rate and the optimal carbon tax are always higher for $\eta = 2/3$ than for $\eta = 2$. However, the difference between the two scenarios decreases significantly under uncertainty as compared to the deterministic case. The optimal carbon tax decreases by 15% to still $72. Appendix A shows that, once more, persistence in the growth shock increases

[11] A reasoning by Nordhaus (2007) suggests that, whenever we decrease η, we should increase the pure rate of time preference in order to keep the overall consumption discount rate fix. We emphasize that this reasoning would be wrong in the current setting. Lowering η implies that we match the observed risk-free rate much better than the standard model. On the other hand, the higher risk aversion parameter explains the higher interest on risky assets, again better than in the standard model. In fact, the empirical literature calibrating the Epstein-Zin model generally finds a lower pure time preference than Nordhaus's (2008) and our $\delta_u = 1.5\%$ along the $\eta = 2/3$ and RAA $= 10$. Given our focus on the effects of uncertainty, however, we decided not to change pure time preference with respect to DICE-2007 in this paper.

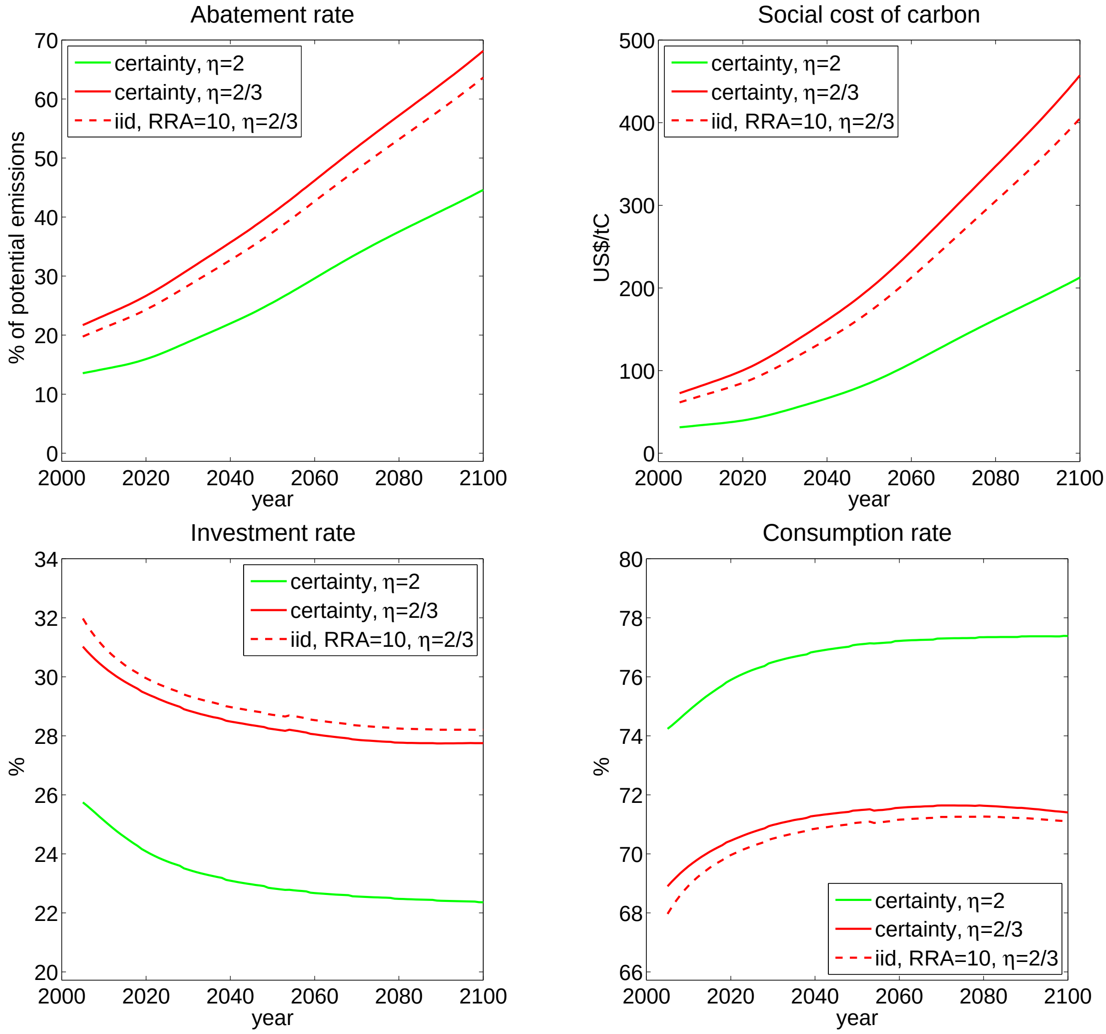

Figure 5 compares the optimal abatement rate, social cost of carbon, investment and consumption under certainty with two different values for the consumption smoothing coefficient, $\eta = 2/3$ and $\eta = 2$, and uncertainty with Epstein-Zin preferences with RRA $= 10$ and $\eta = 2/3$.

the growth uncertainty effect, further reducing optimal policy.

Figure 6 analyzes the dependence of the uncertainty effect on the propensity to smooth consumption over time. We find that growth uncertainty has no effect on abatement for $\eta = 1.1$. At higher levels of η uncertainty increases abatement, at lower levels abatement is higher under certainty. For investment and consumption, we observe no such shift. Uncertainty always increases the investment rate and decreases the consumption rate. These effects slightly decrease in η, implying that the uncertainty effect on investment is slightly lower when the investment rate is already high because of the low consumption smoothing preference.

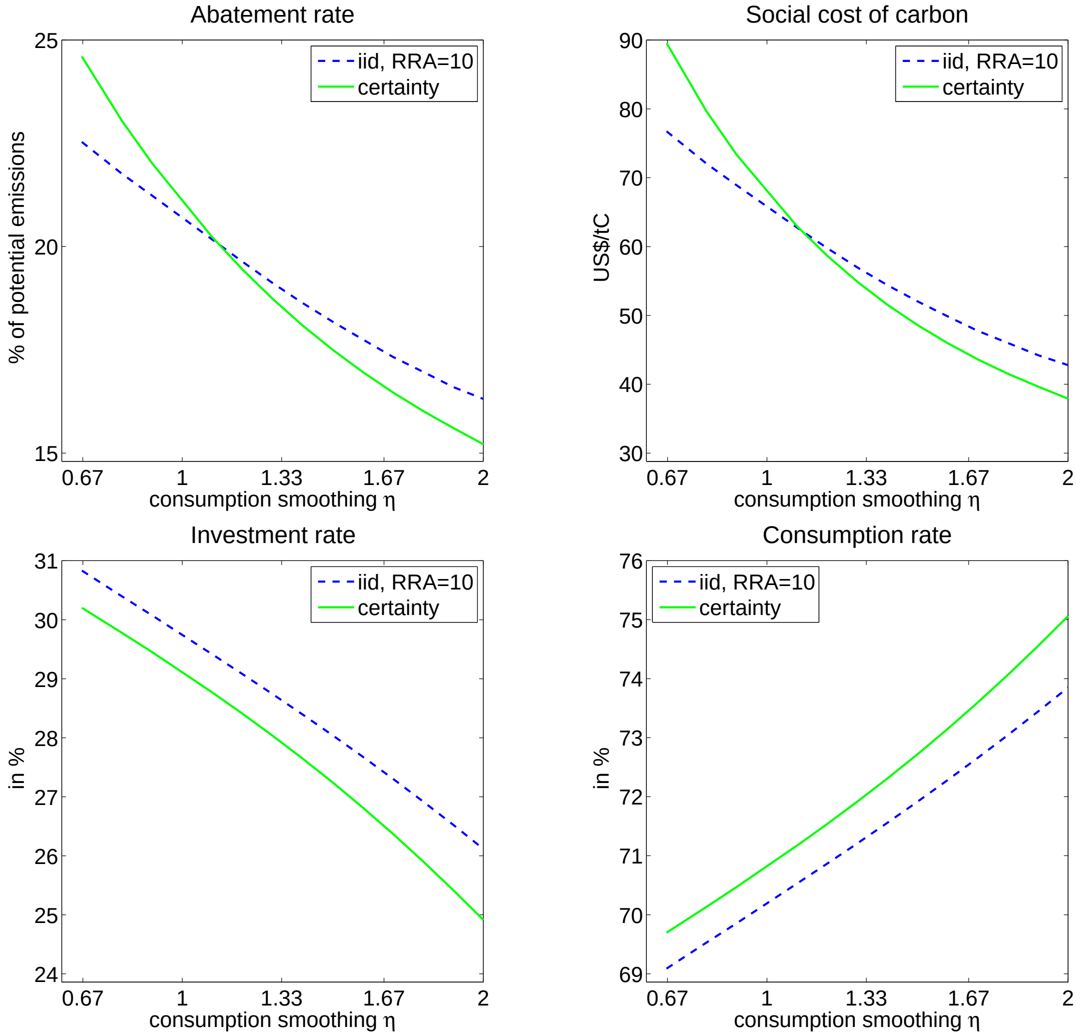

Figure 6 shows the social cost over carbon, abatement, investment and consumption rates under certainty and uncertainty (iid shock with RRA = 10) in the year 2012 for different levels of consumption smoothing η.

4 Discussion

This section explains the precautionary savings observed under Epstein-Zin-Weil preferences and the relation between the consumption smoothing parameter and the uncertainty effect on optimal abatement. In gaining analytic insight, we also discuss how not just different parametric choices, but also the isoelastic functional forms adopted from the original DICE model influence our results. Our discussion simplifies the formulas by expressing consumption and capital in per effective labor units.[12]

[12]The appendix contains the equations of motion in per effective labor terms, i.e. $c_t = \frac{C_t}{A_t^{det} L_t}$, $k_t = \frac{K_t}{A_t^{det} L_t}$ and $y_t = \frac{Y_t}{A_t^{det} L_t}$ (see Appendix D, and for the according Bellman equation Appendix B). Appendix C contains the detailed derivation of the formulas discussed in this section.

4.1 Precautionary savings effect

The previous section observed a large increase in optimal investment and, thus, savings under uncertainty. This precautionary savings effect builds on a prudence effect that is somewhat similar to the discussion of Kimball's (1990) in the standard model. However, the strong precautionary effect that we seek to explain only appears with Epstein-Zin-Weil preferences. It depends on a measure for the difference between (Arrow-Pratt) risk aversion and the propensity to smooth consumption over time. In our setting, this measure is $f(z) = ((1-\eta)z)^{\frac{1-\mathrm{RRA}}{1-\eta}}, z \in \mathbb{R}, (1-\eta)z > 0$. This measure is concave whenever a decision maker is more Arrow Pratt risk averse (measured by RRA) than averse to intertemporal fluctuations (measured by η).[13] Traeger (2010) axiomatically characterizes concavity of f as a measure of intertemporal risk aversion (see section 3.1).

The first order condition for consumption optimization implies

$$u'(c_t) \propto \Pi_t \; \mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial k_{t+1}} \tag{5}$$

Our proportionality drops exogenous terms that do not change under uncertainty or with the preference specification (including the discount factor).[14] Under certainty, and in the entangled standard model, $\Pi_t = P_t = 1$, and the first order condition states that the marginal utility from consumption is proportional to the value derived from investing one more unit into the future capital stock. For notational convenience, we suppress the states of the value function (and its derivatives). Under uncertainty, the technology state of the value function is random and with it the marginal value of capital. The term

$$\Pi_t = \frac{\mathbf{E}_t f'(V_{t+1})}{f'(f^{-1}\mathbf{E}_t f(V_{t+1}))}$$

is known as a prudence term. It is larger than unity if and only if the function f satisfies $-\frac{f'''}{f''} > -\frac{f''}{f'}$ Traeger (2011).[15] An equivalent condition is that the measure of absolute intertemporal risk aversion $-\frac{f''}{f'}$ falls in welfare. It is well known and easily verified that this condition is satisfied in our power function setting unless $\eta = \mathrm{RRA}$ and the power is unity. Therefore, in the disentangled model, the prudence term Π_t always increases the right hand side of equation (5), requiring a higher marginal utility in equilibrium and, thus, lower consumption. The intuition is straight forward: Under decreasing absolute risk aversion, the decision maker has an incentive to save for a higher future welfare so she suffers less from uncertainty.

[13] More precisely it is concave and increasing or convex and decreasing both of which imply that the corresponding aversion measure $-\frac{f''}{f'} > 0$.

[14] Equation (9) in Appendix C contains the proportionality factor, given by the pure time preference and a term relating to expected growth in the effectivity of labor.

[15] The same condition is derived by Kimball (1990) for positive precautionary premia in the standard expected utility model using Arrow Pratt risk aversion.

The term

$$P_t = \frac{f'(V_{t+1})}{\mathbf{E}_t f'(V_{t+1})}$$

is known as a pessimism term. P_t is a normalized weight that fluctuates together with the technology shock. It carries the name pessimism term because, for a concave function f, high welfare realizations translate into a low weight P_t, while low welfare realizations result into a high weight. The decision maker effectively biases the probabilities of bad outcomes upwards. Whether this term increases the expected value in equation (5) depends on whether high realizations of the value function are accompanied by high realizations of the marginal value of capital $\frac{\partial V_{t+1}}{\partial k_{t+1}}$. A theoretical prediction of this relation in complex models is generally hard to derive. In our climate change application, we find that the value function and its derivative vary in opposite directions (or are anticomonotonic in the technology shock). This finding holds for high and low degrees of intertemporal substitutability. As a consequence, the realizations with high capital value receive higher weights, and the pessimism bias increases the right hand side of equation (5). Therefore, also the pessimism term increases savings under uncertainty. Observe that the pessimism effect relies on intertemporal risk aversion directly, rather than on the higher order derivative.

4.2 Abatement effect

The first order condition for optimal abatement implies

$$\Lambda'(\mu_t) \propto \frac{\mathbf{E}_t \; P_t \; \left(-\frac{\partial V_{t+1}}{\partial M_{t+1}}\right)}{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial k_{t+1}}} \; . \tag{6}$$

We dropped a positive proportionality constant that only depends on the period t state of the system and is not affected by uncertainty or changes in the preference specification. The $\Lambda'(\mu_t)$ on the left hand side of equation (6) denotes the marginal expenditure as a fraction of total production to abate one more unit of emissions, measured as fraction of business as usual emissions. The equation states that, in the optimum, this marginal abatement cost is proportional to the expected damage from a ton of carbon emitted and inversely proportional to the value of a production unit. By Jensen's inequality, the convexity of the two different derivatives of the marginal value function in the technology stock decides about the sign of the uncertainty effect. A high convexity of marginal damages in the technology level would increase optimal marginal abatement cost and, thus, the abatement rate. In contrast, a more convex marginal value of capital implies that uncertainty increases the value of capital and, thus, reduces the fraction of production spent on abatement. The main message conveyed by equation (6) is that the abatement rate is determined by the difference in the effect that growth uncertainty has on marginal damages versus the marginal value of capital. This observation fleshes out the difference to the precautionary savings decision characterized in equation (5) and points out why the uncertainty effect on abatement is more ambiguous. Moreover, the prudence term, which was a major driving force of precautionary savings, is absent in equation (6) – it affects produced capital and the value of a lower carbon stock equally and cancels. Similarly, the pessimism effect appears in both expectations. As a consequence, the uncertainty effect only depends on the difference of the pessimism effect when acting on expectation over the carbon damages versus the capital value. Equation (6) gives insight into the trade-offs at stake and how they depend on uncertainty. However, the equation does not enable us to gain an intuition why the sign of the uncertainty effect depends on the propensity to smooth consumption over time.

In Appendix C we eliminate the value function from the right hand side of equation (6) and obtain

$$\Lambda'(\mu_t) \propto \mathbf{E}_t^* \sum_{\tau=t}^{\infty} \left\{ \prod_{j=t}^{\tau} \beta_j \Pi_j P_j \right\} \frac{u'(c_{\tau+1})}{u'(c_t)} \left(-\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \right) \frac{\partial M_{\tau+1}}{\partial M_{t+1}} \ . \tag{7}$$

The expectation operator $\mathbf{E}_t^*$ takes expectations over all possible future sequences $A_{t+1}, A_{t+2}, ...$ (as opposed to just A_{t+1}), conditional on A_t. Equation (7) shows that the optimal marginal abatement costs are a discounted sum over all future marginal damages resulting from increasing current emissions. The term $\prod_{j=t}^{\tau} \beta_j \Pi_j P_j$ is a prudence and pessimism adjusted discount factor for damages in period τ.[16] The term $\frac{\partial M_{\tau+1}}{\partial M_{t+1}}$ accounts for the decay of carbon over time.[17] The additional atmospheric carbon in period τ induces a production loss proportional to $\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}$, which is evaluated in present value consumption units using the conversion factor $\frac{u'(c_{\tau+1})}{u'(c_t)}$.

We use equation (7) to analyze the effects of a technology shock. Ceteris paribus, this persistent growth shock increases production in all subsequent periods. An increase in future production has two dominant effects. First, an increase in $y_{\tau+1}$ proportionally increases the marginal damage from a ton of carbon.[18] Second, an increase in $y_{\tau+1}$ increases period $\tau+1$ consumption, reducing the present value of the damage. We assume for the time being that the consumption rate is constant.[19] Then $\frac{u'(c_{\tau+1})}{u'(c_t)} \propto y_t^{-\eta}$: the marginal present value loss of a consumption unit decreases at the power $-\eta$ as the future grows richer. Together, these two implications of production growth imply that every period's contribution to the present-consumption-equivalent cost of an additional emission unit is proportional to $y_t^{1-\eta}$. Growth uncertainty increases abatement if and only if this function is convex. The convexity condion is

[16] The discount factor $\beta_t = \exp[-\delta_u + g_{A,\tau}(1-\eta) + g_{L,\tau}]$ discounts utility from period $t+1$ to period t units. It picks up a time index to adjust for labor and expected technology growth. These exogenous rates shift from the utility function into the discount factor as we normalize our Bellman equation in Appendix B. As a consequence, the marginal utility from consumption per effective unit of labor does not have to be weighted with the growing population.

[17] This decay is governed by $\frac{\partial M_{\tau+1}}{\partial M_{t+1}} = \prod_{j=t+1}^{\tau} \left[(1-\delta_{M_t,t}) + \frac{\partial \delta_{M,t}}{\partial M_t}(M_t - M_{pre}) \right]$.

[18] $\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}|_{k,M,T} = -g(M,T,t)\, y_t$ where $g(M,T,t)$ depends on the fixed states of the climate system only.

[19] Golosov et al. (2011) spell out conditions that imply a constant consumption rate in a closely related setting. Apart from the Cobb-Douglas production that we also adopt, these assumptions include logarithmic utility, a simplified damage formulation, and full depreciation of capital over the time step.

satisfied only if $\eta > 1$: the convexity of the marginal utility function has to be strong enough to imply an overall convexity of the damage terms in production. We emphasize that, in the mechanism at work, $\eta > 1$ characterizes a high convexity of the marginal utility function and, thus, high prudence. DICE's power utility entangles prudence with the aversion to consumption smoothing. Therefore, under power utility, a higher consumption smoothing preference implies a stronger increase of the abatement rate under uncertainty. Figure 6 shows that the effect of uncertainty switches signs close to $\eta = 1$ but not exactly at the point of logarithmic utility. Our argument assumed a constant consumption rate and neglected the prudence and pessimism terms in the adjusted discount factor. In particular, we find that the optimal consumption rate is negatively correlated with the productivity shock: A higher productivity decreases the consumption rate at the expense of investment (even though absolute consumption still increases). Moreover, we find that the consumption rate is mostly concave in A, which can explain the decrease of optimal abatement under uncertainty for $\eta = 1$.

5 Conclusions

Extrapolating current growth into the future implies that climate policy is a redistribution from a relatively poor present generation to far richer future generations. Over time horizons relevant to climate change evaluation these growth assumptions are highly uncertain. We analyze the implication of growth uncertainty on optimal climate policy in a dynamic programming adaptation of the DICE-2007 model. Our shocks on the rate of technological progress make the economy's technology level a random walk. Under standard preferences, stochastic growth has a minor effect on optimal greenhouse gas abatement and the optimal carbon tax. However, the standard economic model also shows a similar insensitivity to risk in financial markets, giving rise to the equity premium puzzle (too low a risk premium) and the risk-free rate puzzle (too high a discount rate). To evaluate climate change under uncertainty, we emphasize the importance of getting the discount rate and the risk premium right. We therefore follow an approach suggested in the finance literature resolving the puzzles by disentangling risk aversion from a decision maker's propensity to smooth consumption over time. The resulting model satisfies the same rationality constraints as the standard discounted expected utility model, including time consistency.

Increasing relative risk aversion to the degrees measured in finance significantly increases optimal mitigation policies under uncertainty. In particular, the present optimal carbon tax increases by over 20% under an iid shock to over \$40. Introducing a moderate persistence to the shock doubles the uncertainty effect on both the carbon tax and the optimal abatement rate implying an optimal carbon tax just above \$50. The empirical findings in the corresponding finance literature also suggest a lower aversion to intertemporal consumption smoothing than in DICE-2007. Such a reduction in aversion turns the effect of uncertainty on optimal climate policy on its head. Abatement now decreases in response to uncertainty. The optimal carbon tax falls by 15%. However, it does so from a much higher level: a lower aversion to intertemporal substitution decreases the consumption discount rate and increases optimal mitigation. Thus, the fully disentangled model still results in a highest abatement rate, and an optimal carbon tax of $72, but not because of uncertainty. It is merely a consequence of better capturing the low risk-free discount rate.

A different aspect of the optimal policy under growth uncertainty is to increase investment into produced capital. Here, the sign is unambiguous and the magnitude significant. We explain this effect by a measure of intertemporal risk aversion that captures the difference between a decision maker's Arrow Pratt risk aversion and his aversion to intertemporal consumption change. We consider both a prudence and a pessimism effect. In particular, our isoelastic preferences imply that a decision maker is prudent and reduces the welfare impact of growth uncertainty by means of precautionary savings. These prudence and pessimism terms are less influential in determining the optimal abatement policy. The latter depends on the ratio of the marginal value of an emission reduction and the marginal value of capital. Prudence and pessimism terms affect both in similar ways, balancing each other. We therefore derive an analytic formula for the marginal abatement cost that directly relates the optimal abatement rate to a (prudence and pessimism adjusted) discounted sum of future damages. We show that the convexity of marginal utility together with the linear response of marginal damages to production explain the sign switch of the uncertainty effect. Only if marginal utility is sufficiently convex the utility-prudence effect dominates the damage related growth response, and the optimal policy increases abatement under uncertainty. Thus, not consumption smoothing, but the entangled higher order prudence, i.e. change of the smoothing preference with consumption level, determines the sign of the uncertainty effect.

Our paper employs observed preference specifications that are fully rational. In the context of climate change, future wealth is the wealth consumed by future generations not currently alive. Instead of employing observed preferences, we could argue for the use of normative evaluation criteria. Then, equality of generation over time would most likely play a prominent role. Our simulation, as well as straight forward social discounting arguments, show how a low intergenerational substitutability over time (high aversion) implies higher emissions under certainty. In this scenario, uncertainty aversion has again a strong enhancing effect on optimal mitigation efforts.

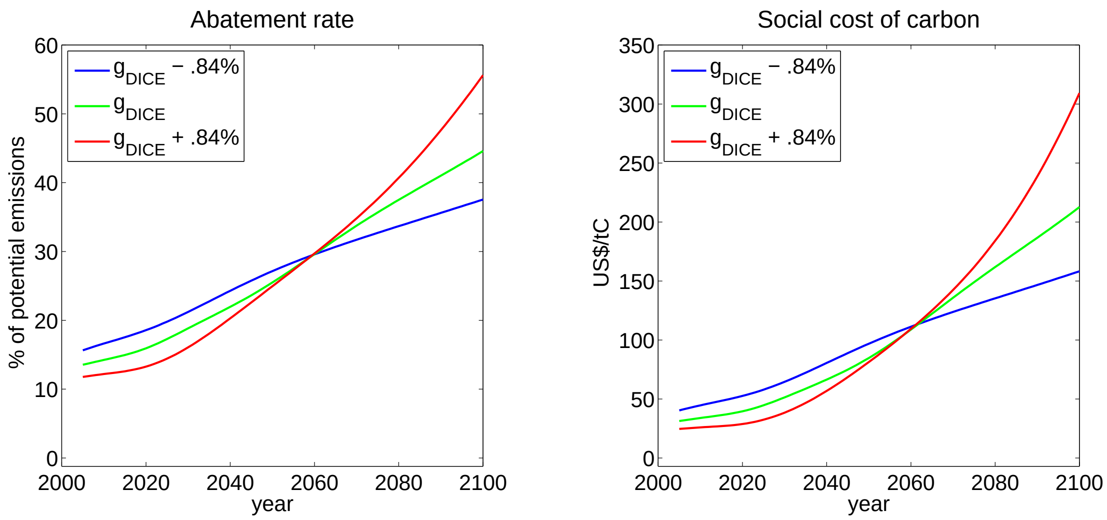

Figure 7 compares the optimal abatement rate and social cost of carbon under certainty with the DICE growth rate $g_{DICE,t}$, a high and a low growth rate ($g_{DICE,t} \pm 0.84\%$).

Appendix

A Further results

Figure 7 shows the impact of varying the growth rate in a deterministic environment. The three growth rates represented correspond to the original DICE-2007 growth rate, a 0.84 percent decrease, and a 0.84 increase at all times. The left panel in Figure 7 shows the optimal abatement rate and the right panel shows the optimal social cost of carbon (SCC). The differences in the three time paths reflects the importance of growth for the timing and level of abatement. The higher the deterministic growth rate, the lower the initial CO_2 abatement: Wealth is taken from rich future generations and transferred to the relatively poorer current generations by depreciating environmental capital. In the lowest growth scenario the optimal policy never reaches full abatement (not shown). With relatively high growth, abatement increases steeply, is between 12 and 13 percent higher at the end of the century, and reaches full abatement more than 50 years earlier as compared to the DICE-2007 baseline. Observe that the deterministic growth rate changes all imply a non-monotonic change of the abatement rate with respect to the original deterministic DICE-2007 baseline. In contrast, our uncertainty simulation all change the optimal climate policy into a single direction, increasing abatement and SCC for $\eta = 2$ and decreasing abatement and SCC for $\eta = 2/3$.

Figure 8 shows the impact of technological uncertainty on optimal abatement and the social cost of carbon for 1000 random path realizations of the stochastic $\tilde{z}$. At the end of the century, there is a 5 per cent chance that the abatement rate is lower than 38 per cent or higher than 65 per cent, the median being 48 per cent. The social cost of carbon is above 400 or below 160 US Dollars per ton of carbon, with the expected value being 250 Dollars.

Figure 9 shows that a probability weighted averaging of deterministic runs has

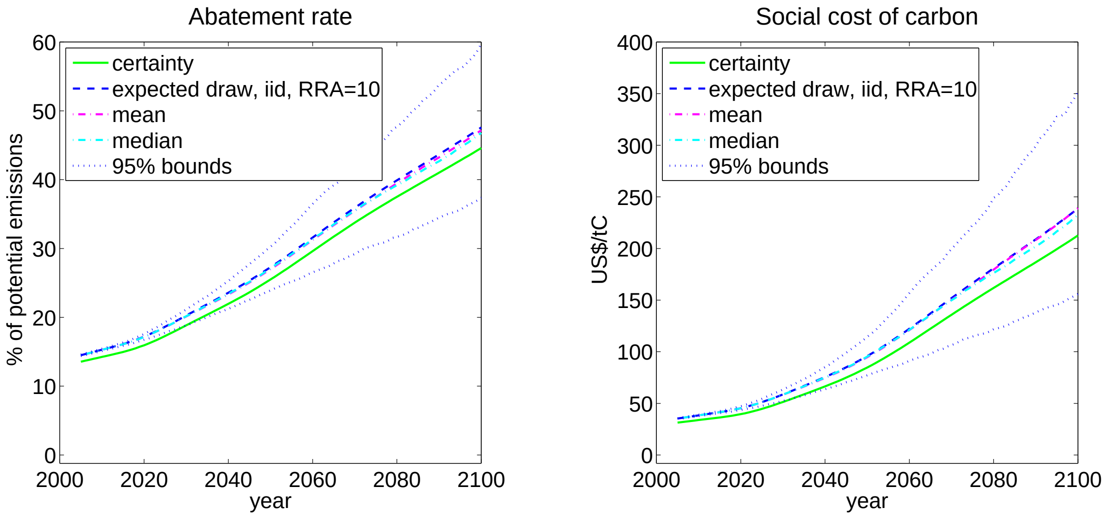

Figure 8 shows the mean, the median, the expected draw and the 95 % confidence bounds for 1000 random paths for abatement and the social cost of carbon.

almost no effect on optimal policy.[20] Such probabilistic averaging, or Monte-Carlo analysis, of deterministic runs is sometimes performed as a first approximation to modeling uncertainty.

Figure 10 shows that persistence in the growth shock also increases the negative effect of uncertainty on mitigation in the setting with a low propensity to smooth consumption over time, where $\eta = 2/3$ (and RRA = 10). Numerically the case of $\eta = 2/3$ is harder than the case where $\eta = 2$ because the parameter choice effectively reduces the contraction of the Bellman equation (8). Thus, we had to settle for a considerably lower level of uncertainty, still showing how persistence increases the negative effect of uncertainty on mitigation.

Figures 11 and 12 show the result of calibrating our simplified climate module to the original DICE-2007 model. The calibration is the same for both sets of graphs and the differences are similar for both $\eta = 2$ and $\eta = 2/3$. The optimal policies, in particular the optimal abatement policy, and the evolution of the carbon stock match the original DICE paths well. However, in order to calibrate these well, we accept the price that our exogenous transitional feedbacks do not match the heat-capacity related delay equation of temperature as well.

[20]The figure averages five runs corresponding to Gaussian quadrature nodes in a normal distribution over the permanent growth 'shock', where $\sigma(\hat{z}) = g_{A,0}/\sqrt{20}$, $\mathbf{E}[\hat{z}] = -\sigma^2(\hat{z})/2$. The permanent shocks imply major changes to the growth dynamics, including destabilizing the numerical model. Thus, we chose a relatively smaller variance to illustrate the effect of Monte-Carlo averaging as opposed to the one chosen in the truly stochastic model.

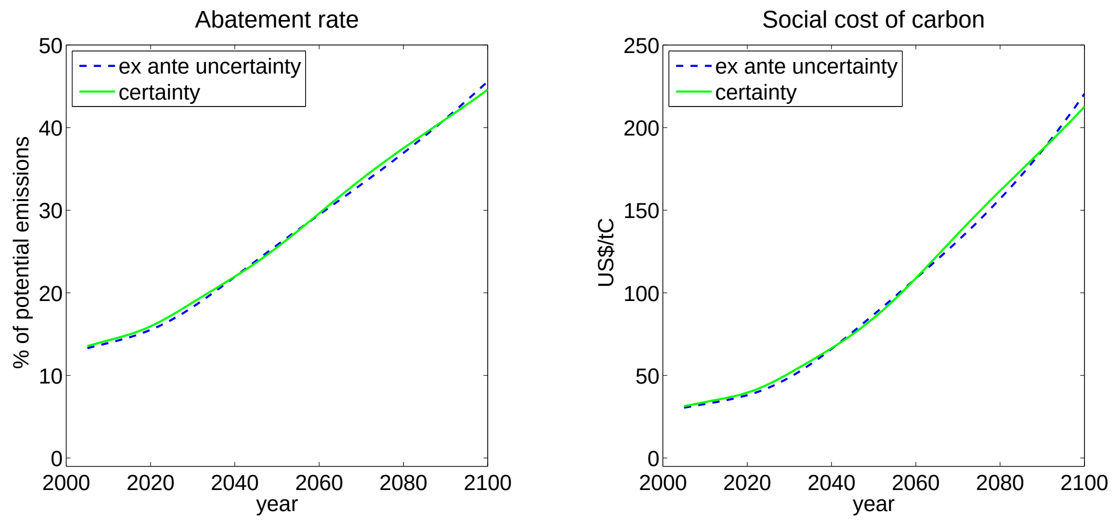

Figure 9 compares the optimal abatement rate and social cost of carbon under certainty and ex ante uncertainty with $\sigma(x) = g_{A,0}/\sqrt{20}$.

B Renormalization of the Bellman equation and numerical implementation

We approximate the value function by the collocation method, employing Chebychev polynomials. We solve the Bellman equation for its fixed point by function iteration. For all models we use seven collocation nodes for each of the state variables captial, carbon dioxide, technology level and the persistent shock. Along the time dimension, we fit the function over ten nodes for the model without, and seven nodes for the model with persistence in the shock. The function iteration is carried out in MATLAB. We utilize the third party solver KNITRO to carry out the optimization and make use of the COMPECON toolbox by Miranda & Fackler (2002) in approximating the value function.

To accomodate the infinite time horizon of our model, we map real time into artificial time by the following transformation:

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

This transformation also concentrates the Chebychev nodes at which we evaluate our Chebychev polynomials in the close future in real time, where most of the exogenously driven changes take place.

Further, we improve the performance of the recursive numerical model significantly by expressing the relevant variables in effective labor terms. Due to the uncertainty in the level of technology, we normalize by the deterministic technology level A^{det}. This is the level of technology under certainty (with all shocks equal zero, $z_t = 0 \ \forall t$)

$$A^{det}_{t+1} = A^{det}_t \exp\left[g_{A,t}\right]$$

Expressing consumption and capital in effective labor terms results in the definitions $c_t = \frac{C_t}{A^{det}_t L_t}$ and $k_t = \frac{K_t}{A^{det}_t L_t}$. Moreover, we define $a_t = \frac{A_t}{A^{det}_t}$. The normalized

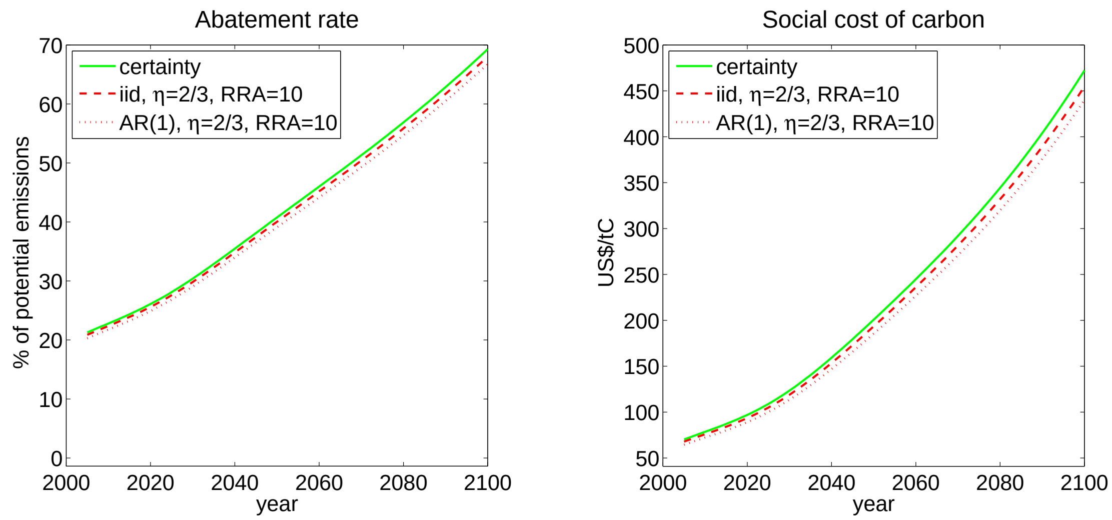

Figure 10 compares the optimal abatement rate and the social cost of carbon under certainty, iid uncertainty and persistent uncertainty with persistence $\zeta = 0.5$ for RRA $= 10$ and $\eta = 2/3$.

productivity one period ahead is then defined as

$$\tilde{a}_{t+1} = \frac{\tilde{A}_{t+1}}{A_{t+1}^{det}} = \frac{\exp\left[\tilde{g}_{A,t}\right] A_t}{\exp\left[g_{A,t}\right] A_t^{det}} = \exp[\tilde{z}] a_t \ .$$

Using all of those new variables we can transform the Bellman equation (4):

$$\frac{V(A_t^{det} L_t k_t, M_t, A_t^{det} a_t, t, d_t)}{(A_t^{det})^{1-\eta} L_t} = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \frac{\exp[-\delta_u + g_{A,t}\ (1-\eta) + g_{L,t}]}{1-\eta} \times \left(\mathbf{E} \left[(1-\eta) \frac{V(A_{t+1}^{det} L_{t+1} k_{t+1}, M_{t+1}, A_{t+1}^{det} \tilde{a}_{t+1}, t+1, \tilde{d}_{t+1})}{(A_{t+1}^{det})^{\rho} L_{t+1}} \right]^{\frac{1-\text{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\text{RRA}}} .$$

Using in addtion artificial time τ, we define the new value function

$$V^*(k_\tau, M_\tau, a_\tau, \tau, d_\tau) = \left. \frac{V(K_t, M_t, a_t A_t^{det}, t, d_t)}{\left(A_t^{det}\right)^{1-\eta} L_t} \right|_{K_t = k_t A_t^{det} L_t,\, t = -\frac{\ln[1-\tau]}{\iota}} ,$$

which leads to the new Bellman equation

$$V^*(k_\tau, M_\tau, a_\tau, \tau, d_\tau) = \max_{c_\tau, \mu_\tau} \frac{c_\tau^{1-\eta}}{1-\eta} + \frac{\beta_\tau}{1-\eta} \times \left(\mathbf{E} \left[1 - \eta V^*(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, \tilde{a}_{\tau+\Delta\tau}, \tau + \Delta\tau, \tilde{d}_{\tau+\Delta\tau}) \right]^{\frac{1-\text{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\text{RRA}}} . \tag{8}$$

When expressing capital and consumption in effective units of labor, we need to adjust the discount factor $\beta_\tau = \exp[-\delta_u + g_{A,\tau}(1-\eta) + g_{L,\tau}]$ by labor and productivity growth. In 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 (see equation 12). This switch of variables turns the constraints on the optimization problem linear.

We recover the original value function from

$$V(K_t, M_t, A_t, t, d_t) = V^*\left(\frac{K_t}{A_\tau^{det} L_\tau}, M_\tau, \frac{A_\tau}{A_\tau^{det}}, \tau, d_\tau\right) \; \left(A_t^{det}\right)^\rho \; L_\tau\Big|_{\tau=1-\exp[-\iota t]} .$$

The marginal value of a ton of carbon is given by

$$\partial_{M_t} V(K_t, M_t, A_t, t, d_t) = \partial_{M_\tau} V^*(k_\tau, M_\tau, a_\tau, \tau, d_\tau) \; \left(A_\tau^{det}\right)^{1-\eta} \; L_\tau\Big|_{\tau=1-\exp[-\iota t]} ,$$

and similarly the marginal value of an additional unit of consumption is

$$\partial_{K_t} V(K_t, M_t, A_t, t, d_t) = \partial_{k_\tau} V^*(k_\tau, M_\tau, a_\tau, \tau, d_\tau) \; \left(A_\tau^{det}\right)^{1-\eta} \; L_\tau \; \partial_{K_\tau} \frac{K_\tau}{A_\tau^{det} L_\tau}\Bigg|_{\tau=1-\exp[-\iota t]}$$

The social cost of carbon in units of the consumption good (US\$) in current value terms is then given by

$$SCC_t = \frac{\partial_{M_t} V}{\partial_{K_t} V} = \frac{\partial_{M_\tau} V^*}{\partial_{k_\tau} V^*} \; A_\tau^{det} \; L_\tau\Big|_{\tau=1-\exp[-\iota t]} .$$

C Derivation of analytic formulas

The first order condition for consumption optimization in the normalized Bellman equation (8) returns

$$u'(c_t) = \beta_t \underbrace{\exp(-g_{A,t} - g_{L,t})}_{\equiv g_t} \underbrace{\frac{\mathbf{E}_t f'(V_{t+1})}{f'(f^{-1}\mathbf{E}_t f(V_{t+1}))}}_{\equiv \Pi_t} \mathbf{E}_t \underbrace{\frac{f'(V_{t+1})}{\mathbf{E}_t f'(V_{t+1})}}_{\equiv P_t} \frac{\partial V_{t+1}}{\partial k_{t+1}} \tag{9}$$

The first order condition for abatement optimization in the normalized Bellman equation (8) returns

$$\mathbf{E}_t P_t \left[\frac{\partial V_{t+1}}{\partial k_{t+1}} \frac{g_t}{1 + D(T_t)} + \frac{\partial V_{t+1}}{\partial M_{t+1}} \mu'(\Lambda) \sigma_t A_t L_t\right] = 0$$

$$\Rightarrow \Lambda'(\mu_t) = -\underbrace{\frac{\sigma_t A_t L_t}{\frac{g_t}{1+D(T_t)}}}_{\equiv \frac{\alpha(T_t,t)}{g_t}} \frac{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial M_{t+1}}}{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial k_{t+1}}}$$

$$\Rightarrow \Lambda'(\mu_t) = -\alpha(T_t, t) \; \beta_t \Pi_t \frac{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial M_{t+1}}}{u'(c_t)} , \tag{10}$$

where we use equation (9) in the last step. Differentiating the Bellman equation (8) partially with respect to the carbon stock M_t using the envelope theorem returns

$$\begin{aligned}\frac{\partial V_t}{\partial M_t} &= \beta_t \Pi_t \mathbf{E}_t P_t \Bigg[\frac{\partial V_{t+1}}{\partial M_{t+1}} \underbrace{\left[(1-\delta_{M,t}) + \frac{\partial \delta_{M,t}}{\partial M_t}(M_t - M_{pre})\right]}_{\frac{\partial M_{t+1}}{\partial M_t}} + \frac{\partial V_{t+1}}{\partial k_{t+1}} g_t \frac{\partial y_t}{\partial M_t}\Bigg] \\ &= u'(c_t)\frac{\partial y_t}{\partial M_t} + \beta_t \frac{\partial M_{t+1}}{\partial M_t} \Pi_t \mathbf{E}_t P_t \frac{\partial V_{t+1}}{\partial M_{t+1}} \ ,\end{aligned}$$

again using equation (9). Repeated substitution of this relation advancing the time indices by one period implies

$$\begin{aligned}\frac{\partial V_t}{\partial M_t} &= u'(c_t)\frac{\partial y_t}{\partial M_t} + \beta_t \frac{\partial M_{t+1}}{\partial M_t}\Pi_t \mathbf{E}_t P_t \ \ u'(c_{t+1})\frac{\partial y_{t+1}}{\partial M_{t+1}} + \\ &\quad \beta_t \frac{\partial M_{t+1}}{\partial M_t}\Pi_t \mathbf{E}_t P_t \ \ \beta_{t+1}\frac{\partial M_{t+2}}{\partial M_{t+1}}\Pi_{t+1}\mathbf{E}_{t+1}P_{t+1} \ \ \frac{\partial V_{t+2}}{\partial M_{t+2}} \\ &= u'(c_t)\frac{\partial y_t}{\partial M_t} + \sum_{\tau=t}^{\infty}\left\{\prod_{j=t}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j \mathbf{E}_j P_j\right\} \ \ u'(c_{\tau+1})\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \ . \end{aligned} \tag{11}$$

Inserting equation (11) into equation (10) gives us

$$\begin{aligned}\Lambda'(\mu_t) &= -\alpha(T_t,t)\Bigg[\frac{\beta_t \Pi_t \mathbf{E}_t \ P_t \ u'(c_{t+1})\frac{\partial y_{t+1}}{\partial M_{t+1}}}{u'(c_t)} + \\ &\quad \frac{\beta_t \Pi_t \mathbf{E}_t \ P_t \sum_{\tau=t+1}^{\infty}\left\{\prod_{j=t+1}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j \mathbf{E}_j P_j\right\} \ \ u'(c_{\tau+2})\frac{\partial y_{\tau+2}}{\partial M_{\tau+2}}}{u'(c_t)}\Bigg] \\ &= -\frac{\alpha(T_t,t)}{\frac{\partial M_{t+1}}{\partial M_t}}\sum_{\tau=t}^{\infty}\left\{\prod_{j=t}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j \mathbf{E}_j P_j\right\} \ \ \frac{u'(c_{\tau+1})}{u'(c_t)}\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \\ &= -\frac{\alpha(T_t,t)}{\frac{\partial M_{t+1}}{\partial M_t}} \ \mathbf{E}_t^* \ \sum_{\tau=t}^{\infty}\left\{\prod_{j=t}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j P_j\right\} \ \ \frac{u'(c_{\tau+1})}{u'(c_t)}\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \ .\end{aligned}$$

While the expectation operators $\mathbf{E}_t$ take expectations over the realization of A_{t+1} (or the normalized a_{t+1}) conditional on earlier realization of A_t, the operator $\mathbf{E}_t^*$ takes expectations over all possible future sequences $A_{t+1}, A_{t+2}, ...$ conditional on A_t.

D The climate enriched economy model

The following model is largely a reproduction of DICE-2007. The three most notable differences are the annual time step (DICE-2007 features ten year time periods), the infinite time horizon, and the replacement of the carbon sink structure by a decay rate. This simplification is neccessary because each carbon sink would require an own state variable in a recursive framework, which is computationally too costly. All parameters are characterized and quantified in Table D on page 32.

Carbon in the atmosphere accumulates according to

$$M_{t+1} = M_{pre} + (M_t - M_{pre})\,(1 - \delta_M(M,t)) + E_t \ .$$

The stock of CO_2 (M_t) exceeding preindustrial levels (M_{pre}) decays exponentially at the rate $\delta_M(M,t)$. The rate is calibrated to the mimick carbon sink structure in DICE-2007. First we calculate the implicit decay rates for the business as usual (BAU) and the optimal policy scenarios in DICE. For each scenario we then approximate a decay rate function over time by cubic splines. Finally, for any point in time, and for all possible levels of carbon stock, we linearly interpolate between the BAU and the optimal decay functions, using the respective carbon stocks from DICE as weights. Since our aim is not primarily to get the relation between carbon stocks and temperature right but to closely match the optimal policies from DICE, we adjust the decay rate δ_M by a factor of 0.75. This comes at the acceptable cost of temperatures rising slightly too fast and not high enough (see Figures 11 and 12).

The variable E_t characterizes yearly CO_2 emissions, consisting of industrial emissions and emissions from land use change an forestry B_t

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

Emissions from land use change and forestry fall exponentially over time

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

Industrial emissions are proportional to gross production $A_t L_t k_t^\kappa$. They can be reduced by abatement. As in the DICE model, we in addition include an exogenously falling rate of decarbonization of production σ_t

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

The economy accumulates capital according to

$$k_{t+1} = [(1 - \delta_k)\,k_t + y_t - c_t] \exp[-(g_{A,t} + g_{L,t})] \ ,$$

where δ_K denotes the depreciation rate, $y_t = \frac{Y_t}{A_t^{det} L_t}$ denotes 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. Population grows exogenously by

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

Here L_0 denotes the initial and L_∞ the asymptotic population. The parameter g_L^* characterizes the convergence from initial to asymptotic population. Technological progress is exogenously given by equation (1) in section 2.1.

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

$$y_t = \frac{1 - \Lambda(\mu_t)}{1 + D(T_t)} k_t^\kappa$$

where

$$\Lambda(\mu_t) = \Psi_t \mu_t^{a_2} \tag{12}$$

characterizes abatement costs as percent of GDP depending on the emission control rate $\mu_t \in [0, 1]$. The coefficient of the abatement cost function Ψ_t follows

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

with a_0 denoting the initial cost of the backstop, a_1 denoting the ratio of initial over final backstop, and a_2 denoting the cost exponent. The rate g_Ψ describes the convergence from the initial to the final cost of the backstop.

Climate damage as percent of world GDP depends on the temperature difference T_t of current to preindustrial temperatures and is characterized by

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

Nordhaus (2008) estimates $b_1 = 0.0028$ and $b_2 = 2$, implying a quadratic damage function with a loss of 0.28% of global GDP at a 1 degree Celsius warming.

Temperature change T_t relative to pre-industrial levels is determined by a measure for the CO_2 equivalent greenhouse gas increase Φ_t, climate sensitivity s, and transient feedback adjustments χ_t

$$T_t = s \ \Phi_t \ \chi_t \ .$$

In detail, climate sensitivity is

$$s = \frac{\lambda_1 \lambda_2 \ln 2}{1 - f_{eql}} \ ,$$

the measure of equivalent CO_2 increase is

$$\Phi_t = \frac{\ln(M_t / M_{pre}) + EF_t / \lambda_1}{\ln 2} \ ,$$

where exogenous forcing EF_t from non-CO2 greenhouse gases, aerosols and other processes is assumed to follow the process

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

Note that it starts out slightly negatively. Our transient feedback adjustment is given by

$$\chi_t = \frac{1 - f_{eql}}{1 - (f_{eql} + f_t)} \ .$$

The parameter f_{eql} is a summary measure of time-invariant feedback prosesses, i.e. the difference between temperature at time t and the equilibrium temperature for a given carbon stock. The function $f_t = f_t(M, t)$ is the transient feedback, capturing mainly heat uptake by the oceans. It is calibrated to match the implied transient feedback in DICE, in a procedure analogous to the decay rate calibration above. Figures 11 and 12 compare the performance of our model to the original DICE 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.

Aguiar, M. & Gopinath, G. (2007), 'Emerging market business cylces: The cycle is the trend', *Journal of Political Economy* **115**(1), 69–102.

Anthoff, D. & Tol, R. (2010), 'Climate policy under fat-tailed risk: An application of fund', *ESRI Working Paper* (348).

Azfar, O. (1999), 'Rationalizing hyperbolic discounting', *Journal of Economic Behavior and Organization* **38**, 245–252.

Baker, E. & Shittu, E. (2008), 'Uncertainty and endogenous technical change in climate policy models', *Energy Economics* **30**.

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

Bansal, R. & Yaron, A. (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *The Journal of Finance* **59**(4), 1481–509.

Chen, X., Favilukis, J. & Ludvigson, S. C. (2011), 'An estimation of economic models with recursive preferences', *NBER* (17130).

Cian, E. D. & Tavoni, M. (2011), 'Mitigation portfolio and policy instruments when hedging against climate policy and technology uncertainty', *Environmental Modeling & Assessment* **17**, 123–136.

Crost, B. & Traeger, C. P. (2010), Risk and aversion in the integrated assessment of climate change, Department of Agricultural & Resource Economics, UC Berkeley, Working Paper Series 1354288, Department of Agricultural & Resource Economics, UC Berkeley.

Dietz, S. (2009), 'High impact, low probability? an empirical analysis of risk in the economics of climate change', *Climatic Change* **108**(3), 519–541.

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

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

Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2011), 'Optimal taxes on fossil fuel in general equilibrium', *NBER Working Paper* .

Ha-Duong, M. & Treich, N. (2004), 'Risk aversion, intergenerational equity and climate change', *Environmental and Resource Economics* **28**(2), 195–207.

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

Hope, C. (2006), 'The marginal impact of co2 from page2002: An integrated assessment model incorporating the IPCC's five reasons for concern', *The Integrated Assessment Journal* **6**(1), 19–56.

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

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.

Kimball, M. S. (1990), 'Precautionary saving in the small and in the large', *Econometrica* **58**(1), 53–73.

Kocherlakota, N. R. (1996), 'The equity premium: It's still a puzzle', *Journal of Economic Literature* **34**, 42–71.

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

Lemoine, D. M. & Traeger, C. P. (2010), Tipping points and ambiguity in the integrated assessment of climate change, Department of Agricultural & Resource Economics, UC Berkeley, Working Paper Series 1704668, Department of Agricultural & Resource Economics, UC Berkeley.

Lontzek, T. S., Cai, Y. & Judd, K. L. (2012), 'Tipping points in a dynamic stochastic iam', *RDCEP Working paper series* **No.12-03**.

Miranda, M. J. & Fackler, P. L. (2002), *Applied Computational Economics and Finance*, MIT Press.

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

Newell, Richard, G. & Pizer, W. A. (2003), 'Discounting the distant future: how much do uncertain rates increase valuations?', *Journal of Environmental Economics and Management* **46**, 52–71.

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

Nordhaus, W. D. (2007), 'A review of the Stern review on the economics of climate change', *Journal of Economic Literature* **45**(3), 686–702.

Richels, R. G., Manne, A. S. & Wigley, T. M. (2004), 'Moving beyond concentrations: The challange of limiting temperature change', *Working Paper* .

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

Traeger, C. (2012), 'Why uncertainty matters - discounting under intertemporal risk aversion and ambiguity', *CESIFO WORKING PAPER NO. 3727* .

Traeger, C. P. (2011), 'Discounting and confidence', *CUDARE working paper 1117* .

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

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

Webster, M., Santen, N. & Parpas, P. (2012), 'An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty', *Computational Management Science* **9**, 339 – 362.

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

Weitzman, M. L. (1998), 'Why the far-distant future should be discounted at its lowest possible rate', *Journal of Environmental Economics and Management* **36**, 201–208.

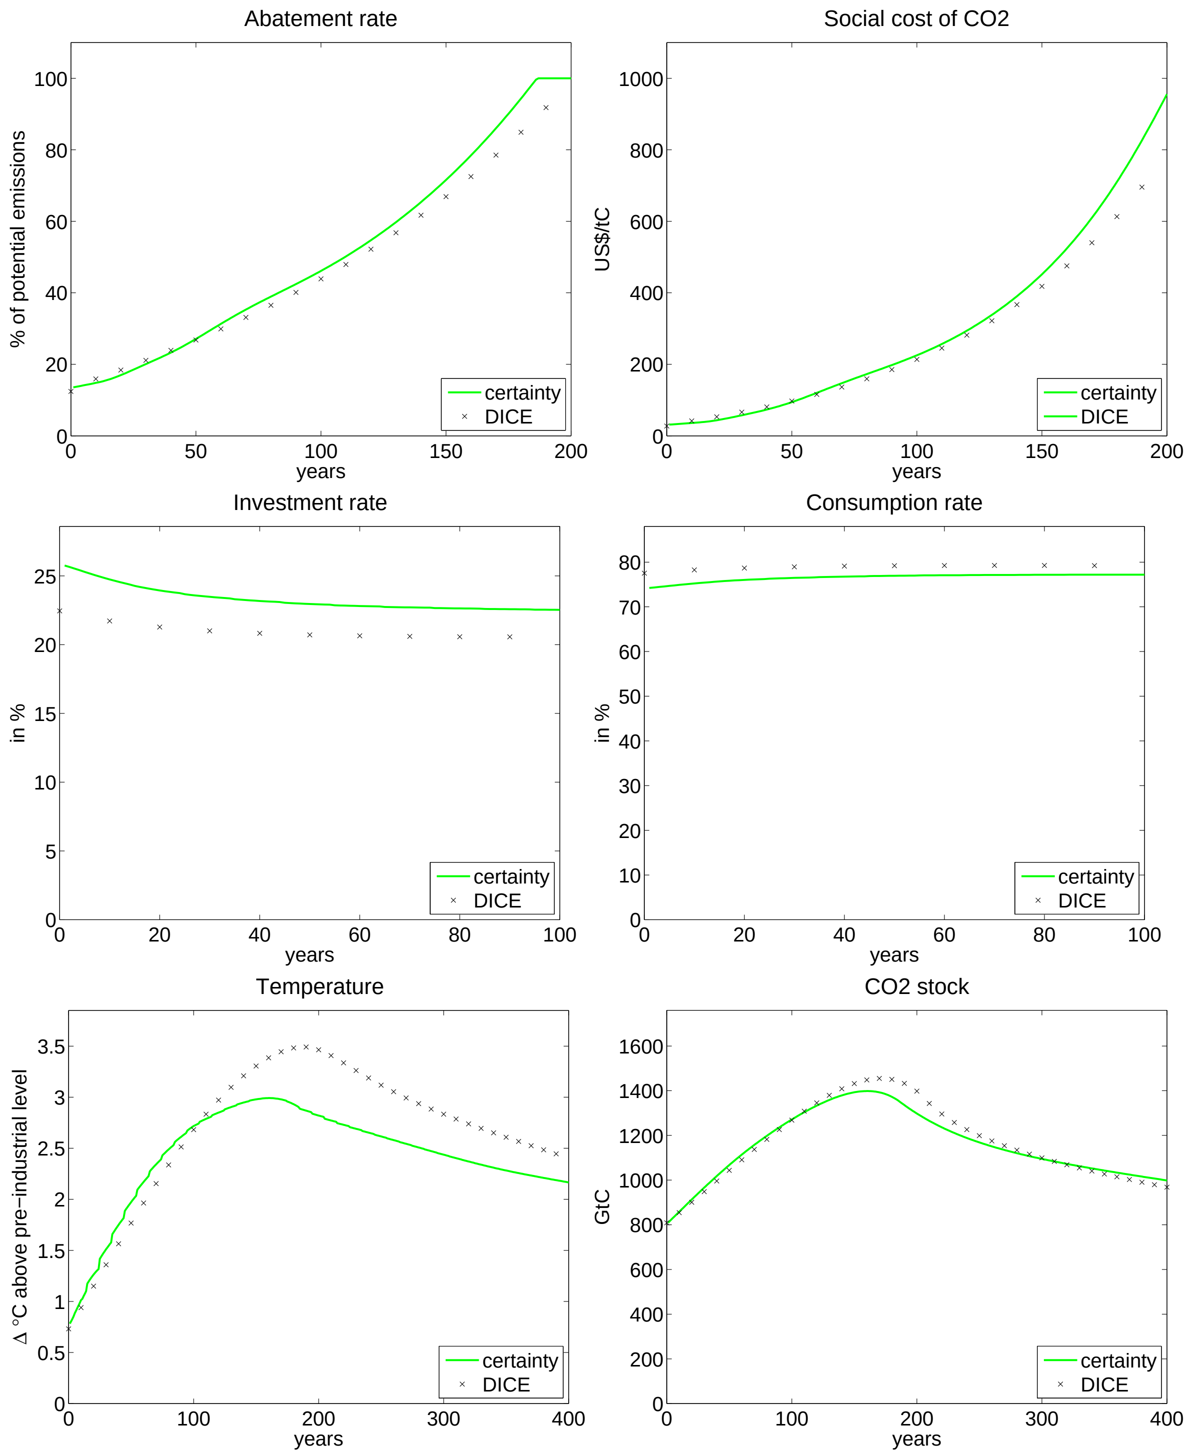

Figure 11: Comparison of the results of our recursive formulation with standard preferences $\eta = 2$ under certainty with the original DICE model results.

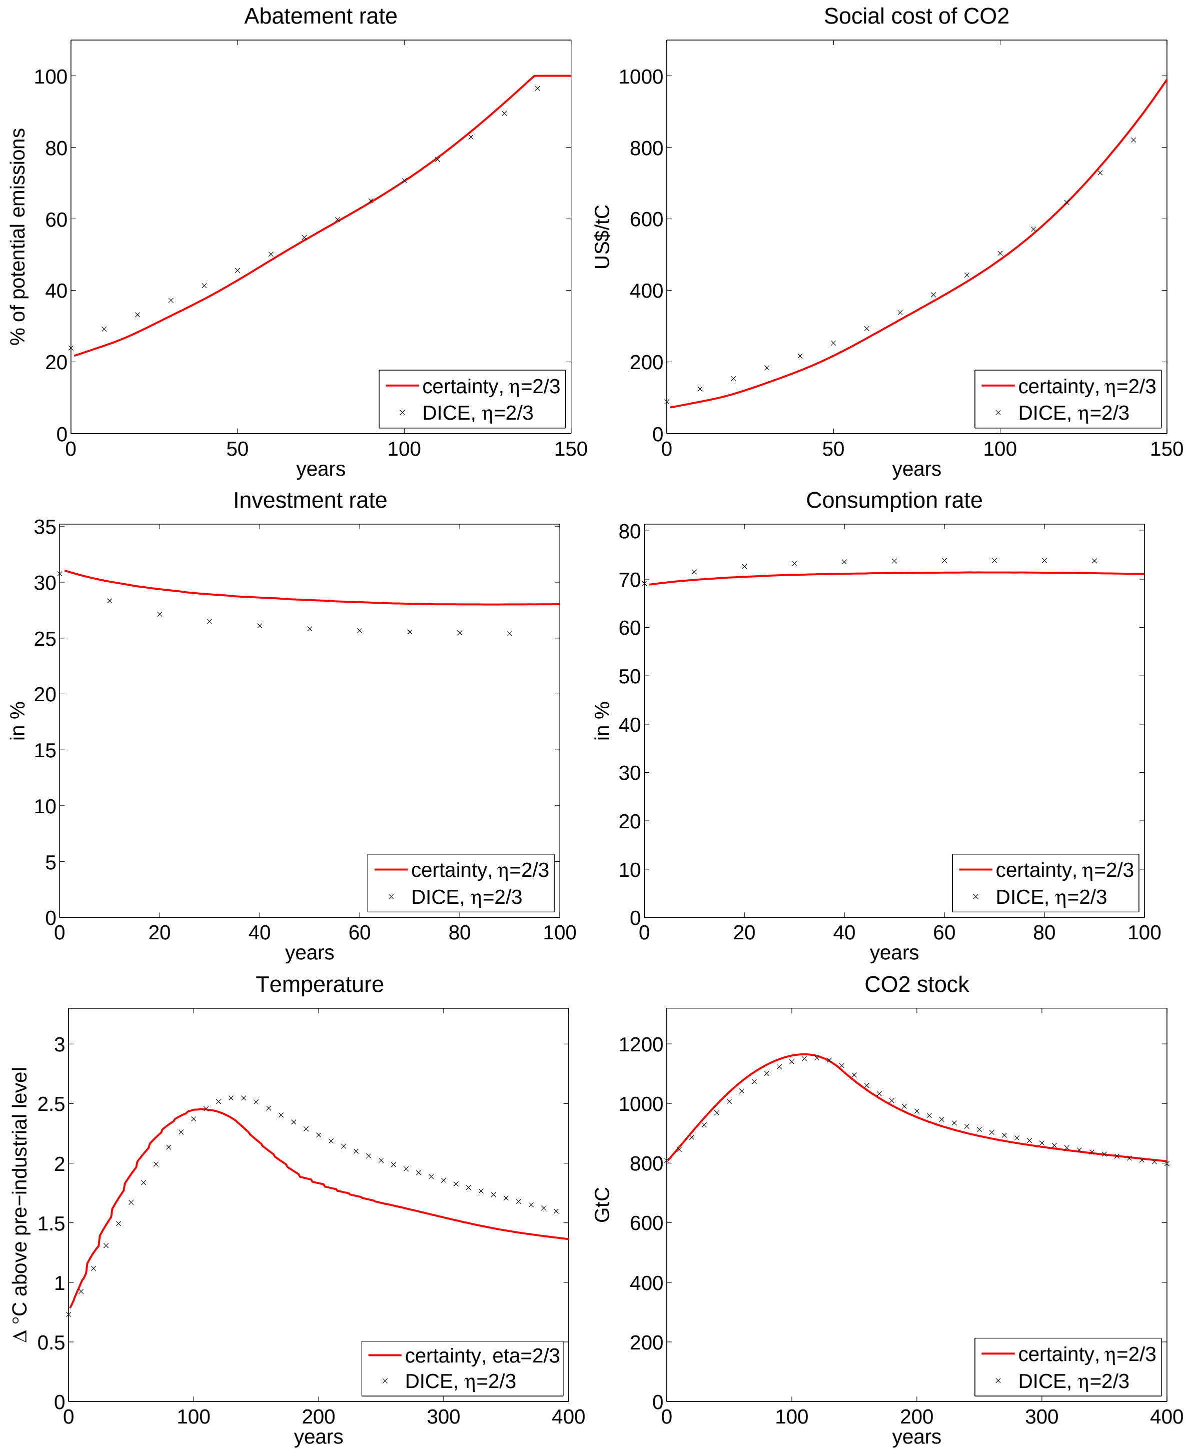

Figure 12: Comparison of the results of our recursive formulation with a low consumption smoothing parameter of $\eta = 2/3$ under certainty with the DICE model results for the same low consumption smoothing parameter.

Table 1 Parameters of the model

Economic Parameters		
η	$\frac{2}{3}, 2$	intertemporal consumption smoothing preference
RRA	$2, 10$	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, preindustiral stock of CO2 in the atmosphere
M_0	808.9	in GtC, stock of atmospheric CO_2 in 2005
$\delta_{M,0}$	1.7%	initial rate of decay of CO2 in atmosphere
$\delta_{M,\infty}$	0.25%	asymptotic rate of decay of CO2 in atmosphere
δ_M^*	3%	rate of convergence to asymptotic decay rate of CO2
B_0	1.1	in GtC, initial CO2 emissions from LUCF
g_B	-1%	growth rate of CO2 emisison from LUCF
s	3.08	climate sensitivity, i.e. equilibrium temperature response to doubling of atmospheric CO_2 concentration with respect to preindustrial concentrations
EF_0	-0.06	external forcing in year 2000
EF_{100}	.3	external forcing in year 2100 and beyond
σ_{forc}	3.2%	warming delay, heat capacity atmosphere
σ_{ocean}	0.7%	warming delay, ocean related
λ_1	5.35	in $W\ m^{-2}$, additional radiative forcing from changing CO2 concentrations 32
λ_2	0.315	in °C $(W\ m^{-2})^{-1}$, temperature change per unit of radiative forcing