Mitigation under Long-Term Growth Uncertainty: Growing Emissions but Outgrowing its Consequences – Sure?[1]

Svenn Jensen† & Christian Traeger§

† Department of Agricultural & Resource Economics, UC Berkeley
207 Giannini Hall #3310, Berkeley, CA 94720-3310, USA; svenn@berkeley.edu

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

Version of August 2013

Abstract: Economic growth over the coming centuries is one of the major determinants of today's optimal greenhouse gas mitigation policy. At the same time, long-run economic growth is highly uncertain. This paper is the first to evaluate optimal mitigation policy under long-term growth uncertainty in a stochastic integrated assessment model of climate change. The sign and magnitude of the impact depend on preference characteristics and on how damages scale with production. We explain the different mechanisms driving optimal mitigation under certain growth, under uncertain technological progress in the discounted expected utility model, and under uncertain technological progress in a more comprehensive asset pricing model based on Epstein-Zin-Weil preferences. In the latter framework, the dominating uncertainty impact has the opposite sign of a deterministic growth impact; the sign switch results from an endogenous pessimism weighting. All of our numeric scenarios use a DICE based assessment model and find a higher optimal carbon tax than the deterministic DICE base case calibration.

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, Matti Liski, Leo Simon, Tony Smith and an anonymous referee for helpful comments.

1 Introduction

The paper analyzes, numerically and analytically, the consequences of growth uncertainty for climate policy in a stochastic integrated assessment. Economic growth is one of the most important determinants of optimal greenhouse mitigation. First, ceteris paribus, economic growth increases carbon emissions in the future. These emissions accumulate in the atmosphere, cause damages for centuries, and aggravate the marginal damage resulting from of a ton of carbon released today. Second, economic growth increases expected future wealth, which makes current generations less willing to sacrifice consumption today to fight future damages from climatic change. Deterministic integrated assessment models usually treat climate change policy as a sure redistribution from the poor to the rich. Nordhaus's (2008) widespread integrated assessment model DICE illustrates this point: even in the absence of any climate policy, the generations living in the year 2100 are five times richer than those living today. Thus, a high propensity to smooth consumption over time (or generations) implies a low optimal carbon tax.

At the same time, economic growth predictions for the far future are highly uncertain. Today's optimal investment into mitigation depends on the co-evolution and interaction of the climate and the economy over centuries. The macroeconomic models underlying integrated assessment extrapolate growth from the past century into the long-run future. We currently cannot foresee whether the explosive growth of the last century will last. Nor can we exclude that growth accelerates further. The current paper is the first to analyze the consequences of long-run growth uncertainty on optimal mitigation policy in an integrated (stochastic dynamic programming) model of the climate and the economy. We quantify the implications of uncertainty using a stochastic version of the DICE model (Nordhaus 2008, Traeger 2012), while our analytic discussion of the underlying mechanisms uses more general functional forms.

The integrated assessment literature predominantly addresses uncertainty by averaging deterministic Monte-Carlo runs (Richels et al. 2004, Hope 2006, Nordhaus 2008, Dietz 2009, Anthoff et al. 2009, Anthoff & Tol 2009, Ackerman et al. 2010, Interagency Working Group on Social Cost of Carbon 2010, Pycroft et al. 2011, Kopp et al. 2012). This first approximation to stochastic analysis can be misleading when deriving optimal policies (Crost & Traeger 2013). Nordhaus (2008, 2011) addresses growth uncertainty along a business as usual trajectory employing the Monte-Carlo approach. His analysis suggests that growth uncertainty reduces the optimal mitigation effort. In contrast, our stochastic analysis shows that growth uncertainty slightly increases optimal mitigation when using DICE's standard preference structure. The critical assumption in the DICE model driving this effect is prudence and the relation between damages and production. Prudence measures the decreases of risk aversion in income. Under growth uncertainty, the prudent agent always invests more into produced capital. However, investment in emission reductions (climate capital) only increases if prudence dominates the production elasticity of damages.

The standard discounted expected utility model entangles intertemporal consumption smoothing with risk aversion, thus, assuming a form of intertemporal risk neutrality (Traeger 2009, 2013). Epstein-Zin-Weil preferences disentangle intertemporal consumption smoothing from risk aversion. First, such a disentanglement allows us to separately identify the effects of consumption smoothing and risk aversion. Second, Epstein-Zin-Weil preferences improve the calibration to observed discount rates and risk premia (Vissing-Jørgensen & Attanasio 2003, Bansal & Yaron 2004, Bansal et al. 2010, Nakamura et al. 2010, Chen et al. 2011, Bansal et al. 2012), both of which are highly relevant for an assessment of climate policy under uncertainty. We find that in the more comprehensive asset pricing framework, growth uncertainty has a much stronger impact on the optimal carbon tax. Using the same propensity to smooth consumption over time as in the DICE base case, growth uncertainty increases the carbon tax between 20% and 45%, depending on whether shocks are iid or (moderately) persistent. However, the corresponding asset pricing literature finds that disentangled estimates of consumption smoothing are much lower than assumed in the DICE model. Using this lower estimate, growth uncertainty decreases the optimal carbon tax between 15% and 30%. An endogenous pessimism weight drives these much larger uncertainty corrections in the more comprehensive asset pricing framework. The reduction of the consumption smoothing preference also reduces the hesitation of current generations to clean up the environment for richer future generations. The uncertainty effect in the Epstein-Zin-Weil setting has the opposite sign of this deterministic consumption smoothing effect. While partially offsetting each other, the deterministic consumption smoothing effect dominates, and all of our quantitative simulations give rise to an optimal carbon tax above the DICE base case level.

Growth uncertainty is also the formal underpinning of Weitzman's (2001) work on falling discount rates for climate change evaluation, which influenced the British Treasury to adopt falling discount rates for the long-term impact of its legislation (Treasury 2003). Weitzman's reasoning assumes permanent uncertainty over the growth rate without learning. Our model implements moderately persistent shocks on the growth rate with rational, anticipated learning. Two recent papers investigate short term deviations from the growth trend in a climate change context. Heutel (2011) studies how optimal unilateral carbon emission reductions by the US vary with the business cycle. Fischer & Springborn (2011) compare labor and output responses to business cycles under different climate policy instruments such as taxes and quotas. Both papers assume stationarity and use (log-)linearizations around the steady state to solve their models. In contrast, we analyze long-term growth uncertainty and solve the full non-linear off-equilibrium dynamics for different isoelastic preference specifications. Our analytic discussion extends beyond the class of isoelastic preferences and particular damage specifications. Traeger (2013) presents a social discounting model for isoelastic preferences suggesting that growth uncertainty has a negligible climate policy impact under entangled expected utility preferences, but a potentially much larger impact under Epstein-Zin-Weil preferences. Baker & Shittu (2008) review the largely theoretical literature on endogenous and often uncertain technological improvements of climate-friendly technologies. Their findings support the use of numerical integrated assessments for technological uncertainty questions, because theoretical results are often ambiguous or dependent on highly stylized models.

Our numeric setting is closest to the projection method approach introduced to integrated assessment by Kelly & Kolstad (1999). To increase numeric efficiency and precision, our model reduces the state space and renormalizes the DICE equations to effective labor units. To employ the Epstein-Zin-Weil coefficients estimated in the long-run risk literature, we employ a one year time step (in an infinite time horizon). These changes to DICE are discussed in detail in Traeger (2012) and briefly summarized in Appendix D. Crost & Traeger's (2010) analysis of damage uncertainty comes closest to our application. Those authors discuss in more detail the consequences of an Epstein-Zin-Weil-based preference calibration for climate change assessment. These calibration effects are distinct (and mostly opposite in sign) from the growth uncertainty implications of Epstein-Zin-Weil preferences discussed in the current paper. Other applications of similar frameworks using the discounted expected utility model include Leach (2007), Kelly & Tan (2013), and Jensen & Traeger (2013) who analyze uncertainty over climate sensitivity, and Lemoine & Traeger (2013f) and Lontzek et al. (2012) who analyze tipping points in the climate system.

Bansal & Yaron (2004) establish the long-run risk model as a prominent explanation of the equity premium and the risk-free rate puzzles, employing persistent shocks on consumption (and dividends) in combination with Epstein-Zin-Weil preferences (Epstein & Zin 1989, Weil 1990, Epstein & Zin 1991). Our stochastic Ramsey-Cass-Koopmans growth model, introduced by Brock & Mirman (1972), is still a "workhorse of modern macroeconomics" (Arouba et al. 2006). In such a production economy, Kaltenbrunner & Lochstoer (2010) confirm that technology shocks generate the long-run consumption uncertainty assumed in Bansal & Yaron (2004). Nakamura et al. (2012) investigate consumption disasters (large, instantaneous and persistent drops in output) using Epstein-Zin-Weil preferences and Nakamura et al. (2012) lend further empirical support to these disentangled preferences in combination with persistent shocks. More generally, the macroeconomic literature answers a vast array of policy questions with related models, including the effects of tax changes on capital income in an open economy (Chatterjee et al. 2004), or the division of risk between capital and labor in a stochastic Romer model (Turnovsky & Smith 2006).

2 Model and Welfare Specification

Our integrated assessment model combines a growing Ramsey-Cass-Koopmans economy with a simple climate model. It is based on the widespread DICE model by Nordhaus (2008) and its stochastic dynamic programming implementation following Kelly & Kolstad (1999) and Traeger (2012). Figure 1 gives a graphical illustration. Production follows a Cobb-Douglas combination of technology (A_t), man-made capital (K_t) and exogenous labor (L_t). Production causes emissions that accumulate in the atmosphere, where they change the Earth's energy balance and cause global warming. Temperatures above the level of 1900 (T_t) reduce future production. To limit future losses the decision-maker reduces emissions. The abatement rate μ is the share of the business as usual emissions that are cut with respect to a laissez-faire regime. We follow DICE in measuring the cost of abatement $\Lambda(\mu)$ as a share of

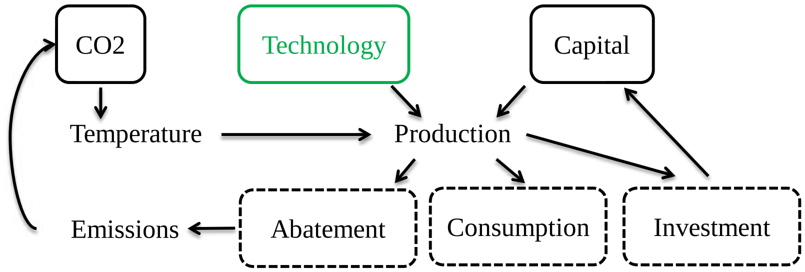

Figure 1 is an abstract representation of the climate-enriched economy model. The control variables consumption and abatement as well as the 'residual' investment are represented by dashed rectangles. The main state variables are depicted by solid rectangles. The green color indicates that the technology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

production. Abatement costs are convex in the abatement rate and they fall exogenously over time. The net production available for consumption and investment into man-made capital is

$$Y_t = \frac{1 - \Lambda_t(\mu_t)}{1 + b_1 T_t^2} (A_t L_t)^{1-\kappa} K_t^{\kappa} \ . \tag{1}$$

The damages $b_1 T^2$ are quadratic in temperatures but (approximately) linear in production. The damage dependence on production will play a major role in characterizing the impact of growth uncertainty on optimal abatement. Our numeric analysis solves for the optimal investment and abatement decisions. Our analytic discussion derives a formula for the optimal marginal abatement cost and, thus, the abatement rate μ. In the following, we discuss in detail uncertain technological progress, welfare, and the Bellman equation. Appendix D gives further model details.

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[2]

$$\tilde{A}_{t+1} = A_t \exp\left[\tilde{g}_{A,t}\right] \quad \text{with} \quad \tilde{g}_{A,t} = g_{A,0} * \exp\left[-\delta_A t\right] + \tilde{z}_t \ . \tag{2}$$

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.[3] We add a stochastic shock $\tilde{z}_t$, which is either identically and independently distributed (iid) or persistent.

[2]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 also leads to balanced growth also in the case of more general production specifications.

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

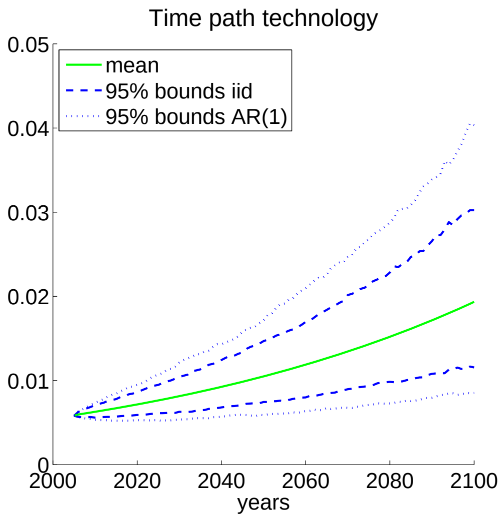

Figure 2 shows the expected draw and the 95% confidence intervals for technology time paths based on 1000 random draws of technology shock $\tilde{z}$ time paths with $\sigma_{\tilde{z}} = 2 * g_{A,0}$. The dotted lines are the confidence intervals for an AR(1), while the dashed lines correspond to iid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

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

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

We base the variance σ_z^2 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 ($\sigma_z = 2 * g_{A,0} \approx 2.6\%$).[4] We fix the mean of the growth shock so that expectations over future technology level coincide with the evolution in the certain scenario (see Appendix E for detailed calculations).[5] 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, Kaltenbrunner & Lochstoer

[4]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.

[5]The technology level in period $t+1$ is lognormally distributed. A mean zero shock of the growth rate would, by Jensen's inequality, imply an increase in the expected next period technology level. Setting $\mathbf{E}[\tilde{z}] = -\sigma^2(\tilde{z})/2$ in every period implies that the $A_{t+\tau}$ expectation equals its deterministic part for all $\tau > 0$ (see Appendix E).

2010). Here, we think of the persistent shock as a more fundamental uncertain change affecting technological progress, e.g., times of economic crisis, international conflict, fundamental innovations or the absence thereof. The theoretical literature has established that persistent shocks imply decreasing social discount rates over time (Weitzman 1998, Azfar 1999, Newell & Pizer 2003). We model persistence in form of an AR(1) process

$$\begin{aligned} \tilde{z}_t &= \tilde{x}_t + \tilde{y}_t \quad \text{where} \\ \tilde{x}_t &\sim \mathcal{N}(\mu_x, \sigma_x^2) \quad \text{and} \\ \tilde{y}_t &= \zeta y_{t-1} + \tilde{\epsilon}_t \quad \text{with} \quad \tilde{\epsilon}_t \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)\ . \end{aligned} \tag{3}$$

Choosing the standard deviations $\sigma_x = \sigma_\epsilon = \sqrt{2} * g_{A,0}$ once again 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 again corresponds to the one under certainty, now conditional on $y_t = 0$.[6] 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 interval 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.

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:

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

[6] Appendix E shows that we achieve this equivalence by setting $\mathbf{E}[\tilde{x}] = \mathbf{E}[\tilde{\epsilon}] = -\sigma^2(\tilde{x})/2$.

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 (4). It is a population (L_t) weighted power function of global per capita consumption (C_t/L_t). The parameter η captures two preference characteristics: the desire to smooth consumption over time and Arrow-Pratt relative risk aversion. Following Nordhaus (2008), we set $\eta = 2$. The second term on the right hand side of equation (4) 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 chooses the optimal control variables consumption C_t and abatement rate μ_t (in DICE: emission control rate). Equation (4) 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 against the value of future (man-made) capital. The control μ_t balances immediate consumption against the reductions of future atmospheric carbon (climate capital).

The standard model underlying equation (4) assumes that intertemporal choice over time also determines risk aversion, and the single parameter η 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

$$V(K_t, M_t, A_t, t, d_t) = \max_{C_t, \mu_t} \frac{L_t \left(\frac{C_t}{L_t}\right)^{1-\eta}}{1-\eta} + \frac{\exp[-\delta_u]}{1-\eta} \left(\mathbf{E} \left[(1-\eta) V(K_{t+1}, M_{t+1}, \tilde{A}_{t+1}, t+1, \tilde{d}_{t+1}) \right]^{\frac{1-\mathrm{RRA}}{1-\eta}} \right)^{\frac{1-\eta}{1-\mathrm{RRA}}} . \tag{5}$$

Now, the parameter η captures only the desire to smooth consumption over time (inverse of the intertemporal elasticity of substituion). The parameter RRA depicts the Arrow-Pratt measure of relative risk aversion. In the case $\eta = \mathrm{RRA}$ equation (5) collapses to equation (4). 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), and Bansal et al. (2012). These papers suggest best guesses of $\eta = \frac{2}{3}$ and of relative risk aversion in the proximity of the value RRA = 10. 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 Normalized Bellman Equation and Intertemporal Risk Aversion

The Bellman equations (4) and (5) are not convenient for a numerical implementation for several reasons. First, modeling a random walk without mean reversion is a numerical challenge and the normalized Bellman equation converges significantly better. Second, the support of the non-normalized capital and the absolute technology state grow without bounds, limiting the planning horizon as well as the node density of a numerical implementation. Third, our renormalized Bellman equation takes a more generic form removing population weights, which is convenient for the analytic discussion. Our renormalized technology state variable a captures the deviation from the deterministic technology path in DICE, $A_{t+1}^{det} = A_t^{det} \exp\left[g_{A,0} * \exp\left(-\delta_A t\right)\right]$. We define a as the ratio of the actual and the hypothetical deterministic technology level, $a_t = \frac{A_t}{A_t^{det}}$. Moreover, we express consumption and capital in per effective labor units, $c_t = \frac{C_t}{A_t^{det} L_t}$ and $k_t = \frac{K_t}{A_t^{det} L_t}$. Finally, we map the infinite time horizon on a $[0, 1]$ interval using the transformation $\tau = 1 - \exp[-\iota t]$, which allows us to approximate the value function over the infinite time horizon. With these renormalization we can restate the Bellman equation (5) as

$$V^*(k_\tau, M_\tau, a_\tau, \tau, d_\tau) = \max_{c_\tau, \mu_\tau} u(c_\tau) + \beta_\tau \times \qquad (6)$$

$$f^{-1}\left(\mathbf{E}\left[f\left(V^*(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, \tilde{a}_{\tau+\Delta\tau}, \tau + \Delta\tau, \tilde{d}_{\tau+\Delta\tau})\right)\right]\right),$$

where $u(c) = \frac{c^{1-\eta}}{1-\eta}$ and $f(v) = ((1-\eta)v)^{\frac{1-\mathrm{RRA}}{1-\eta}}, v \in \mathbb{R}, (1-\eta) > 0$. We introduce general functional u and f because they facilitate a more insightful analytic discussion of our findings in section 4. The function f has an interpretation of intertemporal risk aversion that we discuss in the next paragraph, while u is a generic utility function of per capita consumption. Appendix B spells out the detailed derivation and discusses the numeric implementation.

The curvature of the function f in equation (6) captures the difference between Arrow-Pratt risk aversion and aversion to intertemporal change. In the standard discounted expected utility model both coefficients coincide (RRA $= \eta$) and the function f is linear, implying that it does not affect the uncertainty evaluation in the Bellman equation (6). When the Arrow-Pratt coefficient RRA is larger than the consumption smoothing parameter η, as observed in the asset pricing, then the function f is concave. A concave function f implies a risk averse aggregation over the uncertain future value function. Intuitively, concavity of f captures risk aversion with respect to utility gains and losses. More formally, Traeger (2010) characterizes such an aversion to utility gains and losses axiomatically in a choice theoretic context and labels it intertemporal risk aversion. In an intertemporal setting, risk affects evaluation in two different ways. First, it leads to fluctuations in consumption over time. Decision-makers generally dislike fluctuations over time. This dislike is captured by the consumption smoothing parameter η or, more generally, the concavity of the utility function u which is fully determined by deterministic choice. Second, risk makes future outcomes intrinsically uncertain. This aversion to not knowing which future will come true is captured by intertemporal risk aversion, i.e., the concavity of the function f.[7]

3 Numeric Results

We first illustrate the small impact of uncertainty in the entangled standard model. Second, we switch to the disentangled model and increase the coefficient of relative risk aversion to the value suggested in the finance literature. Finally, we analyze the dependence of optimal policy on the propensity to smooth consumption over time. Persistence of the growth shock is discussed alongside the changes to the preference parameters.

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 a random walk in the technology level (iid shock on growth rate, section 2.1). Here, the decision-maker optimizes under uncertainty, but nature happens to still draw expected values at every instance.[8] 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 (2013) explains the smallness of the precautionary effect by pointing

[7]The concavity of the composition of functions $f \circ u$ captures both the aversion to the intertemporal fluctuation caused by risk (measured by u) and the intrinsic aversion (measured by f). This joint uncertainty aversion characterizes the Arrow-Pratt measure (in the isoelastic model RRA).

[8]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 singles out the policy difference arriving only from acknowledging uncertainty when looking ahead. We compare this representation to other possible path representations in Figure 7 in Appendix A.

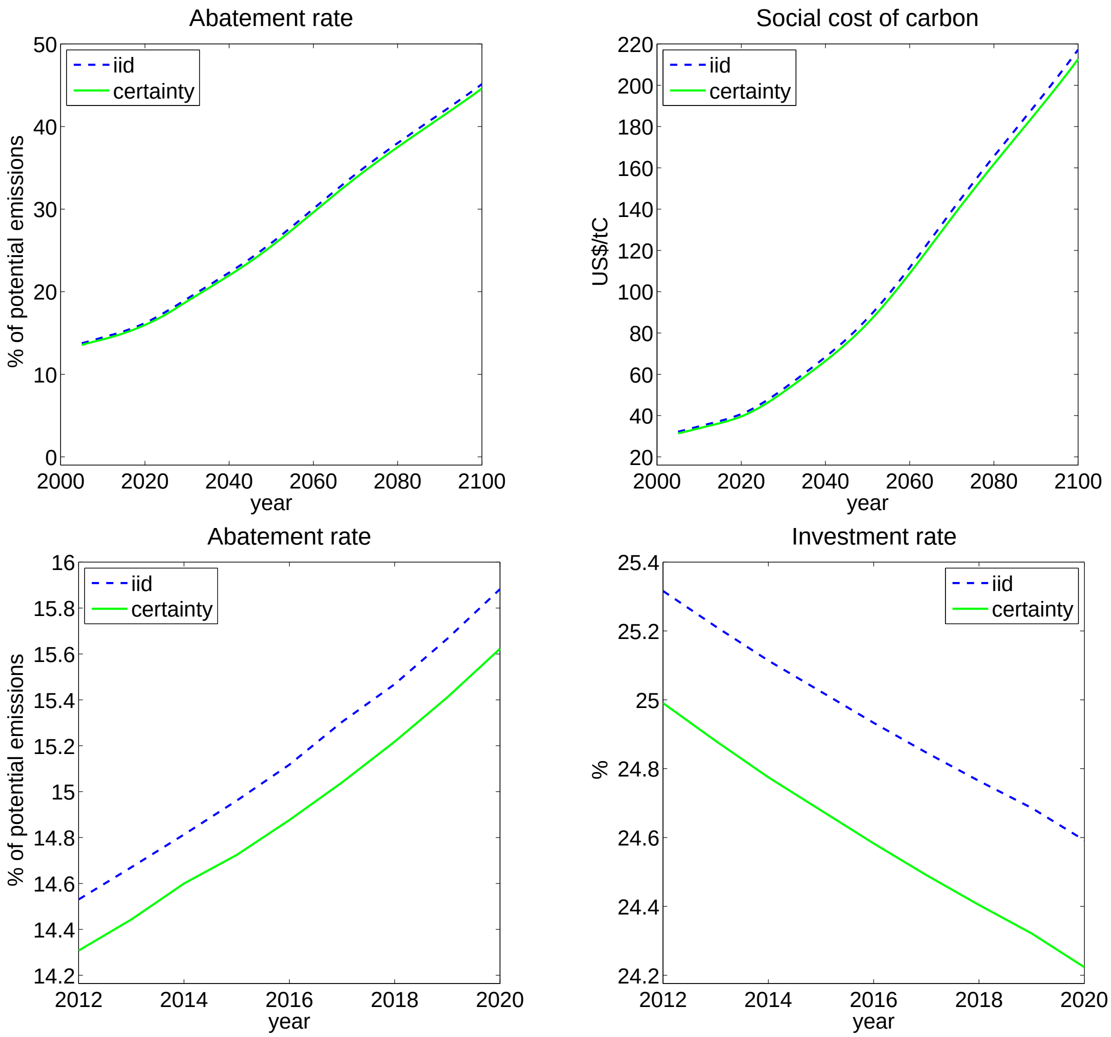

Figure 3 compares the optimal abatement rate, the social cost of carbon, and the investment rate under (iid) uncertainty to their deterministic values (standard preferences, $\eta = 2, \text{RRA} = 2$). The upper panels show the abatement rate and the social cost of carbon for the current century, the lower panels show the abatement rate and the investment rate for the current decade. Uncertainty has a small, positive effect on climate policy and investment.

out that decision-makers with entangled preferences are effectively intertemporal risk neutral (section 2.3).

3.2 Increasing Risk Aversion ($\text{RRA} = 10$)

The standard model of the previous section does not accurately capture risk premia (equity premium puzzle). As we argued in section 2.2, we improve the DICE-2007 calibration to asset markets by employing Epstein-Zin-Weil preferences in the disentangled Bellman equation (5). We increase the relative risk aversion coefficient to $\text{RRA} = 10$, but for now we keep the consumption smoothing parameter at the DICE value ($\eta = 2$).

Figure 4 shows the optimal climate policy keeping aversion to intertemporal sub-

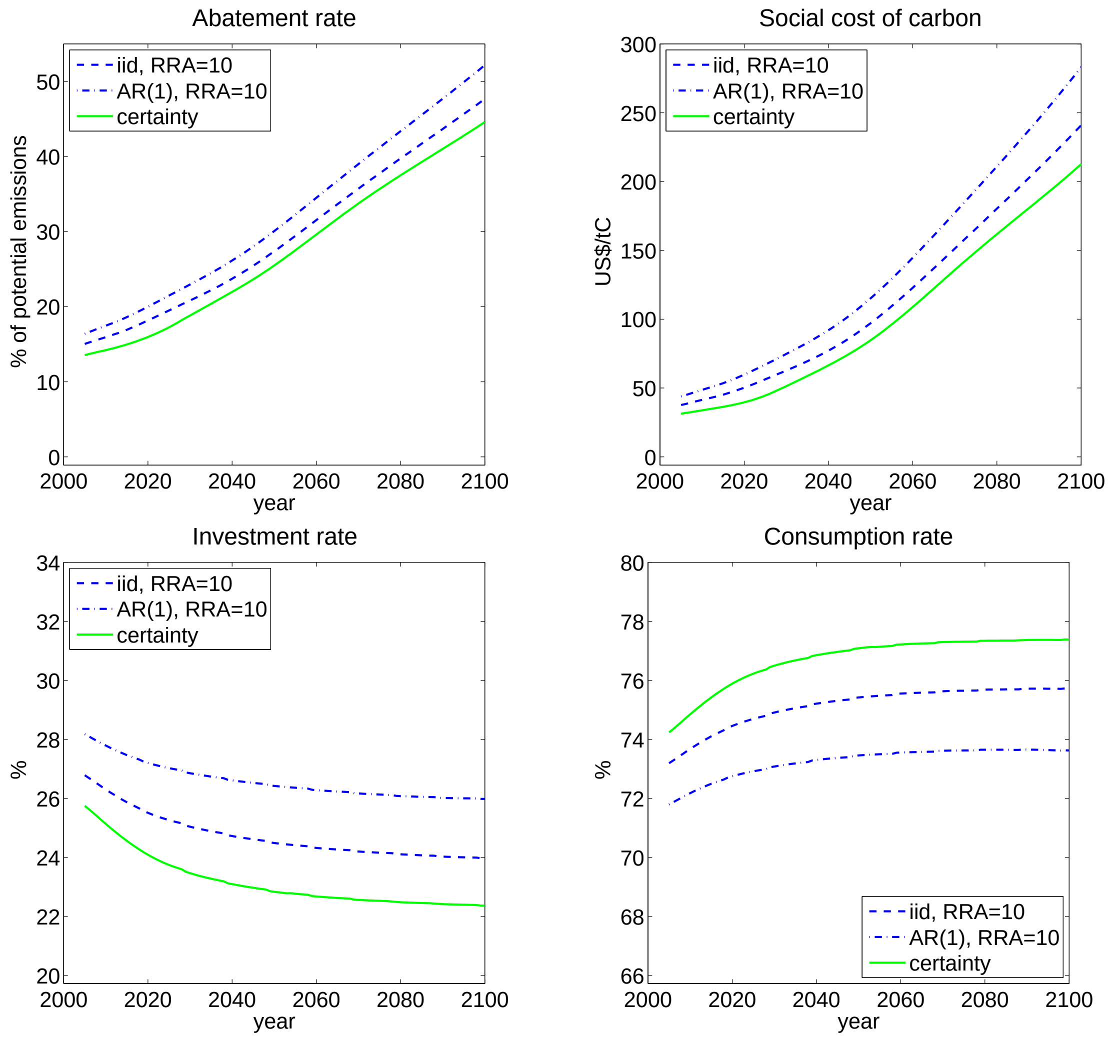

Figure 4 compares abatement, the social cost of carbon, investment rate, and consumption rate for three scenarios: certainty, an iid shock, and a persistent shock. The consumption smoothing coefficient is $\eta = 2$ and the coefficient of relative risk aversion is RRA = 10. Uncertainty increases abatement, the social cost of carbon, and the investment rate. The consumption rate decreases. Persistence magnifies all effects.

stitution at $\eta = 2$ and raising 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 productive 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 (3). 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 reduction of agents' propensity to smooth consumption over time to $\eta = 2/3$ (section 2.2), improving the calibration to the risk-free discount rate.[9] 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 nearly 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 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 $72. The dashed-dotted line shows that, once more, persistence in the growth shock

[9] 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-Weil 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.

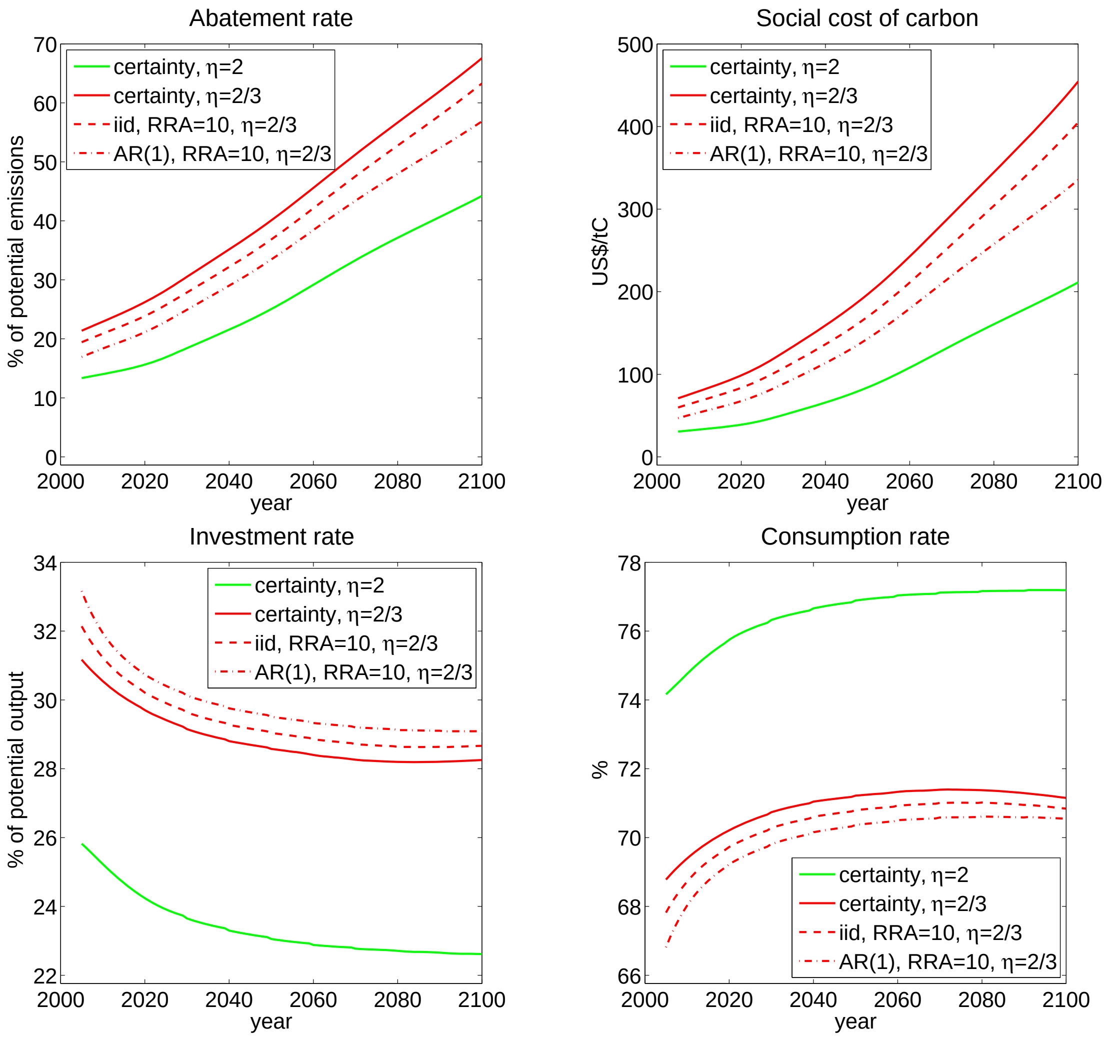

Figure 5 summarizes the results of full preference disentanglement. The solid lines depict certainty scenarios for $\eta = 2/3$ (red/dark) and $\eta = 2$ (green/light). The dashed line represents an iid shock scenario with $\eta = 2/3$, RRA $= 10$, the dashed-dotted line a persistent shock with the same preferences. A weaker desire to smooth consumption over time deterministically increases both the investment rate in man-made capital and the abatement rate (and the carbon tax). Uncertainty (further) increases the investment rate in man-made capital, but reduces abatement.

amplifies the growth uncertainty effect. It reduces optimal mitigation and the social cost of carbon and increases the investment rate.

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.

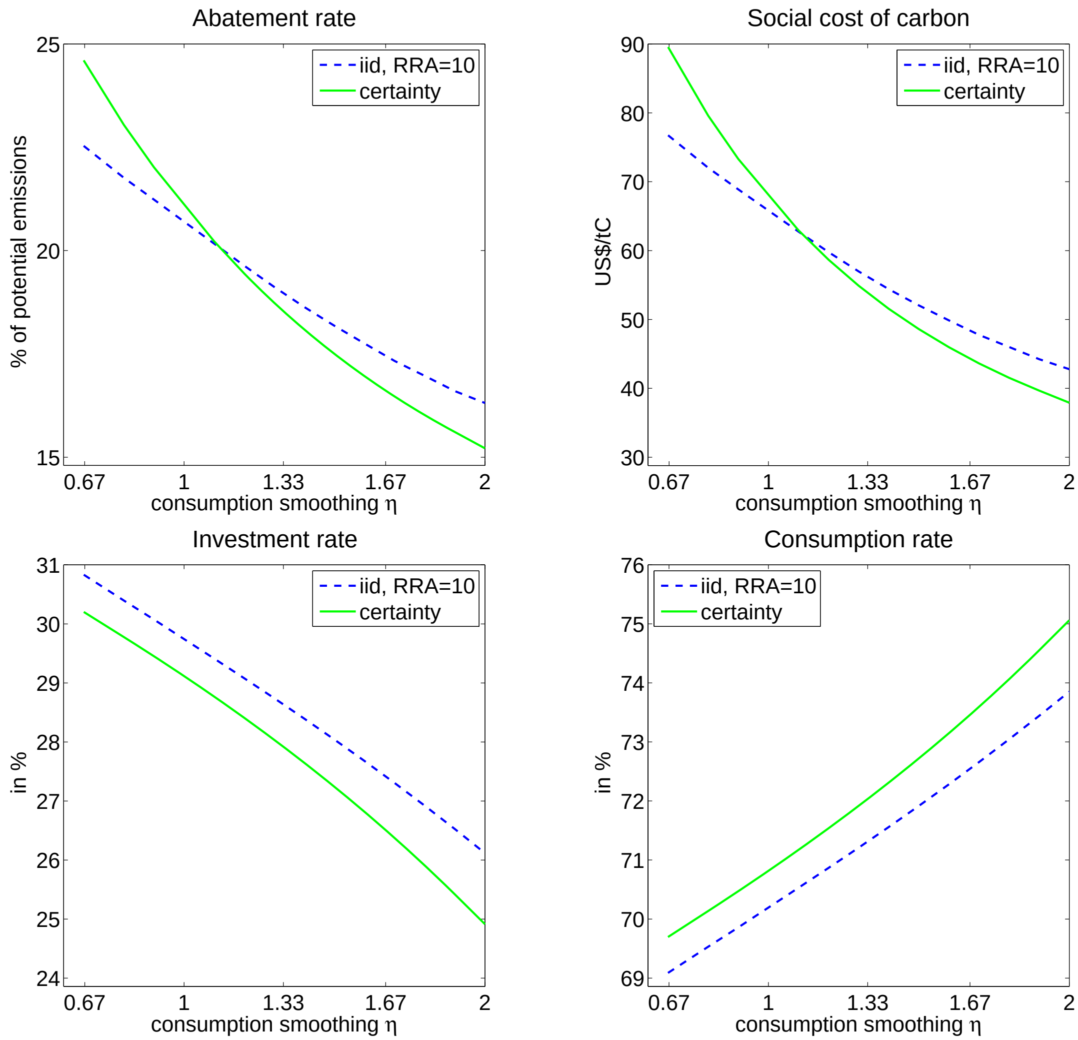

Figure 6 compares the optimal present day controls under certainty and (iid) uncertainty as a function of the consumption smoothing parameter η. Relative risk aversion is RRA $= 10$. For a high value of η, uncertainty decreases the social cost of carbon and (and vice versa). The effect switches sign at $\eta = 1.1$. In contrast, investment in man-made capital always increases under uncertainty.

4 Analytic Discussion

This section develops an analytic understanding of the uncertainty effects observed in the previous section. We start by presenting an analytic formula for optimal mitigation, which paves the ground for our subsequent discussion. Before moving on to a detailed analysis of uncertainty effects, section 4.2 characterizes the deterministic growth trade-off between consumption smoothing and damages sensitivity. The we explain our stylized facts: first, growth risk increases optimal mitigation expenditure for the DICE base specification. Section 4.3 identifies the mechanism and the structural assumptions responsible for this result: the relation between prudence, damage sensitivity, and type of technological progress. Second, the magnitude of the uncertainty effect increases significantly under Epstein-Zin-Weil preferences, and switches sign for a low preference for consumption smoothing. Section 4.4 explains the pessimism weighting mechanism driving these results as well as the observed sign difference as compared to investment in man-made capital. Forth, persistence increases in the growth shock increases the uncertainty effects, which we discuss together with a few other details in section 4.3.2.

4.1 Optimal Mitigation & the Social Cost of Carbon

Mitigating a ton of carbon today decreases the stock of carbon in all future periods. We write the change of atmospheric carbon in period $\tau > t$ as a consequence of a unit reduction of emissions in period t as $-\frac{\partial M_\tau}{\partial E_t}$.[10] The change in period τ carbon level affects period τ output as $\frac{\partial y_\tau}{\partial M_\tau}$. In period τ a unit increase in production increases welfare according to its marginal product $\frac{u'(c_\tau)}{u'(c_t)}$, which we normalize relative to the marginal value of consumption in period t. Thus, under certainty, the social benefit of a unit increase of carbon in period t obtained in period $\tau > t$ is given by the product $-\frac{u'(c_\tau)}{u'(c_t)}\frac{\partial y_\tau}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_t}$. The total benefit accruing from all subsequent periods is the discounted sum of these benefits. In the optimum, these total benefits are proportional to the marginal abatement cost $\Lambda'(\mu_t)$ in period t (see Appendix C):

$$\Lambda'(\mu_t) \propto \mathbf{E}_t^* \sum_{\tau=t}^{\infty} \left\{ \prod_{j=t}^{\tau} \beta_j \Pi_j P_j \right\} \frac{u'(c_{\tau+1})}{u'(c_t)} \left(-\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \right) \frac{\partial M_{\tau+1}}{\partial E_{t+1}} . \tag{8}$$

The proportionality absorbs a positive constant that depends only on the period t state of the system and is not affected by uncertainty or changes in the preference specification. 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.

The first term under the sum $\prod_{j=t}^{\tau} \beta_j \Pi_j P_j$ is a prudence- and pessimism-adjusted discount factor for Epstein-Zin-Weil preferences. The discount factor β_t discounts normalized utility from period $t+1$ to period t units.[11] Π_j is a prudence and P_j is a pessimism adjustment, which arise as a consequence of Epstein-Zin-Weil preferences (section 4.4).

4.2 Growth in the Deterministic Baseline

Ceteris paribus, an positive growth rate shock increases economic production in all subsequent periods. In consequence, future damages are larger, but the marginal damage is valued less by a richer population. Equation (7) allows us to formalize this deterministic trade-off in a growing economy. First, damages are a function of production, the production loss caused by an additional ton of carbon, here $d(y) = -\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}$. In the DICE model, the damage function $d(y)$ is linear in production y and a postive production shock proportionally increases damages (equation 1).[12] Second,

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

[11] This discount factor is $\beta_t = \exp[-\delta_u + g_{A,\tau}(1-\eta) + g_{L,\tau}]$; it includes up a time index to adjust for labor and expected technology growth as a result of normalizing consumption to effective labor units and eliminating the population weight from the Bellman equation. See Appendix B for details.

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

an increase in production $y_{\tau+1}$ increases period $\tau+1$ consumption, reducing the valuation of the marginal damage through the term $u'(c_{\tau+1})$.

For analytic tractability, we henceforth assume a constant consumption rate.[13] The assumption allows us to flesh out the basic mechanisms determining abatement under certainty, under risk, and under Epstein-Zin-Weil preferences. We relax the assumption in section 4.3.3. In equation (7), technological progress (under certainty) affects the terms $u'(c(a))$ and $d(y(a))$. Throughout the discussion section, we will characterize various relations among the relevant variables by means of elasticities. We define the elasticity of damages with respect to production as

$$\mathrm{Dam}_1(d,y) = \frac{d'(y)}{d(y)} y \ .$$

The DICE model assumes $\mathrm{Dam}_1 = 1$. We define aversion to intertemporal substitution as the negative of the consumption elasticity of marginal utility

$$\mathrm{AIS}(u,c) = -\,\mathrm{MU}_1(u,c) = -\frac{u''(c)}{u'(c)} c \ .$$

$\mathrm{AIS}(u,c)$ is the inverse of the intertemporal eslasticity of substitution. For the isoelstic utility function in DICE it is $\mathrm{AIS} = \eta$, and Nordhaus (2008) assumes $\eta = 2$.

In equation (7) we find that technological progress affects the optimal carbon tax (first order) through the terms $u'(c(a))$ and $d(y(a))$, implying that the optimal carbon tax increases under technological progress if $\frac{d}{da} u'(c(a)) d(y(a))$ is positive or, equivalently (see Appendix C)

$$\mathrm{AIS}(u,c) < \mathrm{Dam}_1(d,y) \qquad \Leftrightarrow \qquad \eta < 1 \quad \text{in DICE} \ . \tag{9}$$

Technological growth increases the social cost of carbon and the optimal abatement rate if (and only if) damages are more sensitive to production shocks than the marginal valuation of consumption. This simple insight goes hand in hand with the widely recognized observation that in a growing economy, a lower aversion to intertemporal substitution increases the optimal mitigation level. The latter finding is frequently stated in terms of the social discount rate: a lower consumption smoothing parameter $\mathrm{AIS} = \eta$ reduces the consumption discount rate, increasing the attention paid to long-run climate damages. The right hand-side of equation (9) connects this insight to the straight-forward finding that a higher sensitivity of climate damages to production further increases optimal mitigation.

In DICE, for the base specification where $\eta = 2$, a deterministic increase in the technology level reduces the optimal carbon tax. We show the numeric result in Figure 8 in Appendix A. In the figure, we also observe that the impact of growth on climate policy is non-monotonic over time: towards the end of the century higher

[13] Golosov et al. (2011) spell out conditions that imply a constant consumption rate in a closely related setting. Apart from our Cobb-Douglas production, these assumptions include logarithmic utility, a simplified damage formulation, and full depreciation of capital over the time step. Given our more general setting, the consumption discount rate will not generally be constant and section 4.3.3 briefly discusses the second order effects resulting from a non-constant consumption rate.

growth results in a higher optimal carbon tax. This non-monotonicity is another difference to the uncertainty effects observed in section 3. The rich future generation "makes up" for the lower abatement today, increasing the carbon tax and reaching full abatement significantly earlier. While the characterization of the growth effect in equation (9) also holds in the future, it assumes a given state of the climate system. Comparing the different lines in Figure 8, the carbon stock during the second half of the century is significantly higher for the generations on the high growth path as compared to those on the low growth path. From equation (9) we also observe that with logarithmic utility, the climate policy in a DICE-like model with linear in production damages is independent of the technology (and production) level. This case, where $\eta = 1$ and $\mathrm{Dam}_1(d, y) = 1$, is the setting of Golosov et al. (2011) analytic integrated assessment model.

4.3 The Uncertainty Effect in the Discounted Expected Utility Standard Model

Equation (7) allows us to explain how a growth shock affects optimal abatement. The current section focuses on the trade-off in the entangled standard economic model ($\Pi_j = Pi_j = 1$). First, we analyze the basic intuition using a mean-zero shock on production. Second, we explain how more general forms of technological progress, including our employed labor augmenting progress, modify the characteristics of the uncertainty effect. The analytic discussion in these sections assumes a constant consumption rate. The final section considers the changes arising from an endogenous consumption rate and discusses the effects of persistence in the technology shocks.

4.3.1 Uncertainty over production

To understand how uncertainty affects the social cost of carbon, we now analyze the consequences of a mean-zero shock on production. A non-zero effect of uncertainty arises from asymmetries between the consequences of a positive and a negative growth shock. We continue to denote damages by $d(y) = -\frac{\partial y}{\partial M}$. From Jensen's inequality we know that if the product $u'(c)d(y)$ on the right hand side of (7) is convex, then the optimal carbon tax increases under uncertainty. We therefore introduce the second order normalized moments of the damage function as

$$\mathrm{Dam}_2(d, y) = \frac{d''(y)}{d'(y)} y \ ,$$

which is the elasticity of marginal damage with respect to production. The DICE model assumes that damages are linear in production so that $\mathrm{Dam}_2(d, y) = 0$. Similarly, we define the negative of the second order normalized moment of marginal utility as

$$\mathrm{Prud}(u, c) = -\,\mathrm{MU}_2(u, c) = -\frac{u'''(c)}{u''(c)} c$$

$\text{MU}_2(u, c)$ is Kimball's (1990) measure of relative prudence. It characterizes precautionary savings (=investment) response to income uncertainty. For the isoelstic utility function in DICE we have $\text{Prud} = 1 + \eta = 3$. The positivity of relative prudence explains the increase of investment in produced capital under uncertainty. Finally, note that $\text{RRA}(u, c) = -\text{MU}_1(u, c) = \text{AIS}(u, c)$ ($= \eta = 2$ in DICE) is now also the Arrow-Pratt measure of relative risk aversion. Risk aversion and aversion to intertemporal substitution coincide by assumption in the entangled standard model.

Production uncertainty increases optimal abatement if the marginal value of damages is convex in production, or (see Appendix C)

$$\begin{aligned} \text{Prud}(u, c) &> 2\ \text{Dam}_1(d, y) - \frac{\text{Dam}_1(d, y)}{\text{AIS}(u, c)}\ \text{Dam}_2(d, y) \qquad (10) \\ \Leftrightarrow \quad 1 + \eta &> 2 \ * \ 1 \ - \ 0 \qquad \text{in DICE} . \end{aligned}$$

In contrast to the case of produced capital, prudence not only has to be positive, but it has to dominate the sensitivity of damages to production shocks. In the DICE model, damages are linear in production and utility is isoelastic. Thus, the condition that shocks on production increase optimal abatement simplifies to $\eta > 1$. While η characterizes relative risk aversion and consumption smoothing, we emphasize that the driving force for increased abatement is neither risk aversion nor consumption smoothing dominating unity, but prudence dominating the damage elasticity.

Proportionality of damages to economic production is a ubiquitous assumption in integrated assessment models, but it recently received attention in critical discussions of integrated assessment models (Weitzman 2010). In the extreme case that damages were independent of economic activity, the abatement rate would react similarly to the investment rate in conventional capital. If damages were, e.g., quadratic in the level of production then the damage convexity measure Dam_2 contributes. For a given level of risk aversion, more convex damages reduce the requirements on prudence. The multiplier $\frac{\text{Dam}_1}{\text{AIS}}$ increases this effect if the damage elasticity dominates aversion to intertemporal substitution (or risk aversion), which is the same condition required for an abatement increase under deterministic growth (see equation 9). For isoelastic preferences, prudence ($= 1 + \eta$) and the consumption smoothing preference η are dependent. An increase in η increases both preference measures, increasing both sides of inequality (10). For the example of quadratic-in-production damages, the region where growth uncertainty decreases optimal abatement shifts from $\eta \in [0, 1]$ to the interval $\eta \in [1, 2]$.[14]

4.3.2 Type of technological progress

This subsection relaxes the assumption that production shocks are mean-zero. Our implementation of DICE assumes labor augmenting technological progress, and the shock is designed to keep expected future technology levels the same as under certainty

[14] For quadratic damages in production we find $\text{Dam}_1 = 2$ and $\text{Dam}_2 = 1$ so that equation (10) is positive if and only if $(\eta - 1)(\eta - 2) > 1$ and, thus, the uncertainty effect on abatement is negative if $1 < \eta < 2$.

(see Appendix E). Production is concave in labor and, thus, a shock that keeps the expected technology level constant slightly decreases expected production.[15] Let $\mathrm{y}(a)$ denote the relation between technology and production, and we define once again the normalized moments

$$\mathrm{Tech}_1(\mathrm{y}, a) = \frac{\mathrm{y}'(a)}{\mathrm{y}(a)} a \quad \text{and} \quad \mathrm{Tech}_2(\mathrm{y}, a) = \frac{\mathrm{y}''(a)}{\mathrm{y}'(a)} a \ .$$

Tech_1 is the elasticity of production with respect to technology, and Tech_2 is the elasticity of marginal production with respect to technology (a convexity measure). In the case of our labor augmenting technological progress where $y_{t+1} \propto a_{t+1}^{1-\kappa}$ we find that the linear sensitivity measure is $\mathrm{Tech}_1(\mathrm{y}, a) = 1 - \kappa$, while the convexity measure $\mathrm{Tech}_2(\mathrm{y}, a) = -\kappa$ takes a negative value because production is concave in labor augmenting technology.

Uncertainty increases optimal mitigation under mean-zero shocks on technology a if and only if (see Appendix C)

$$\mathrm{Prud}(u, c) > \mathrm{Dam}_1(d, y) \left[2 - \frac{\mathrm{Dam}_2(d, y)}{\mathrm{AIS}(u, c)} \right] + \frac{\mathrm{Tech}_2(\mathrm{y}, a)}{\mathrm{Tech}_1(\mathrm{y}, a)} \left[1 - \frac{\mathrm{Dam}_1(d, y)}{\mathrm{AIS}(u, c)} \right] \tag{11}$$

As for direct production shocks in equation (10), prudence has to dominate the damage dependence of production (first term on right hand side). In addition, equation (11) accounts for a possible non-linearity in the relation between technology shocks and production (second term on the right hand side). A convexity in the impact of technology a on production y increases expected production. As a consequence, the (prudence) domain for which uncertainty increases abatement becomes larger if $\mathrm{Dam}_1 > \mathrm{AIS}(u, c)$ (reducing the right hand side of inequality 11). This is the same condition that we identified in section 4.2, ensuring that deterministic growth increase abatement. As discussed there, the condition reflects the fact that marginal damages are more sensitive to production changes than is marginal valuation.

In our labor augmenting model with isoelastic utility and linear-in-production damages, equation (11) simplifies to

$$1 + \eta > 2 - \frac{\kappa}{1 - \kappa} \left[1 - \frac{1}{\eta} \right] \qquad \Leftrightarrow \qquad \eta > 1 \ .$$

This criterion coincides with the DICE version of equation (10) for mean-zero production shocks. Thus, under uncertainty $\eta = 2$ indeed increases optimal abatement as observed in Figure 3. For logarithmic utility and linear-in-production damages, uncertainty has no effect on optimal abatement for any type of technological progress. This case includes the uncertainty in Golosov et al.'s (2011) analytic integrated assessment model.

If the damage function was, e.g., quadratic in production, then the range where uncertainty reduces optimal mitigation enlarges from $\eta \in [1, 2]$ under mean-zero total

[15] Observe that except for Figure 7 in the appendix, we depict expected-draw simulations. Thus, along any depicted path nature draws the expected technology level and the actual production level evolves as under certainty, except for changes caused by differences in the optimal policies.

factor productivity shocks (see previous section) to $\eta \in [0.6, 2]$ for mean-zero shocks on labor augmenting technology (assuming DICE's $\kappa = 0.3$).

4.3.3 Endogenous savings and persistent shocks

For the purpose of analytic tractability, the previous section assumed a constant consumption rate. However, Figure 6 shows that the consumption rate decreases under uncertainty. Figure 9 in Appendix A analyzes the consumption control rule and shows that the agent's consumption rate decreases under a positive growth shock and increases under a negative growth shock. Note that absolute consumption slightly increases under a positive growth shock, but much less than overall production. The decrease in the consumption rate for a positive growth shock dampens the change in marginal utility in equation (7). The endogeneity of the consumption rate therefore acts similar to lowering η, reducing the uncertainty effect in the standard model for $\eta = 2$ and increasing the point where the uncertainty effect on mitigation flips from positive to negative. Figure 10 in Appendix A shows the results for the optimal carbon tax when the investment rate is fixed to the deterministic level. While the effect is almost imperceptible for standard preferences, it is small but notable for the case of Epstein-Zin-Weil preferences. Note that overall emissions are still lower because the investment rate is lower (absence of precautionary savings) and, thus, future production is lower.

That persistence increases uncertainty effects is straight forward. Persistent shocks imply the same uncertainty over next period technology as the iid shocks. However, persistence of the shocks implies that more periods in the sum of equation (7) contribute larger effects to the uncertainty correction of abatement. A similar message is conveyed by Figure 2, showing the much larger long-run uncertainty under persistence of shocks.

4.4 Epstein-Zin-Weil Preferences and Intertemporal Risk Aversion

The disentanglement of risk aversion and the propensity to smooth consumption over time permits a more accurate incorporation of risk premia and discount rates in evaluating the climate asset. Our empirical analysis finds a major increase of the uncertainty effects under such a comprehensive preference specification, as well as a sign switch depending on the parameter η. Equation (7) for the social cost of carbon captures the corresponding modifications in terms of the prudence factor Π_j and the pessimism factor P_j. Here, we explain these factors and discuss how they modify optimal climate policy under uncertainty.

4.4.1 Precautionary savings

In Appendix C we show that the first order condition for consumption optimization implies

$$u'(c_t) \propto \Pi_t \ \mathbf{E}_t \ P_t \ \frac{\partial V_{t+1}}{\partial k_{t+1}} \ . \tag{12}$$

The proportionality absorbs exogenous terms that do not change under uncertainty or with the preference specification. 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. An increase on the right hand side of equation (12) increases optimal marginal utility of the last consumption unit and, thus, decreases the consumption level and increases investment.

The prudence term Π_t is defined as

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

For mean-zero shocks over the next period welfare, the prudence term increases the right hand side of equation (12) and, thus, investment under uncertainty if, and only if, absolute intertemporal risk aversion $-\frac{f''}{f'}$ decreases in welfare (Traeger 2011). The prudence label arises from a condition characterizing decreasing absolute risk aversion: $\text{Prud}(f, V) > \text{RRA}(f, V)$, i.e., prudence (of f evaluated at V) dominates (intertemporal) risk aversion. The latter condition is always met for Epstein-Zin-Weil preferences due to their isoelastic form. However, the technology shock does not necessarily produce mean-zero welfare shocks as measured by the value function. For the $\eta = \frac{2}{3}$ scenario we find that the value function is close to linear in a and, thus, the prudence term indeed increases optimal investment. For the $\eta = 2$ scenario, however, we find that the value function is strongly concave in a, biasing down the expected value of V and implying that the prudence term slightly decreases optimal consumption (see Figure 11 in Appendix A). The resulting uncertainty corrections are relatively small and dominated by the pessimism effect discussed in the next paragraph. The difference in curvatures driving the difference in the prudence term between the scenarios results from the stronger decrease of marginal utility in consumption in the $\text{AIS}(u, c) = \eta = 2$ scenario. It outweighs the (intertemporal) risk aversion based prudence effect, keeping the current consumption level slightly higher and investment slightly lower.

The dominating uncertainty impact on consumption and investment operates through the pessimism term defined as

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

P_t is a normalized weight fluctuating with the technology shock. It carries the name pessimism term because, for a concave risk aversion function f, low welfare realizations translate into a high weight P_t, and vice versa. The decision-maker effectively biases the probabilities of bad outcomes upwards. A low realization of technological progress implies a low welfare realization and a high marginal value of capital (in all scenarios). As a consequence, the pessimism bias puts more weight on high realizations of the marginal value of capital and, thus, raises the opportunity cost for consumption (equation 12). Pessimism increases investment and decreases consumption.

4.4.2 Abatement effect

In Appendix C we derive the following first order condition for marginal expenditure on abatement as a fraction of total production:

$$\Lambda'(\mu_t) \propto \frac{\mathbf{E}_t \; P_t \; \left(-\frac{\partial V_{t+1}}{\partial M_{t+1}}\right)}{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial k_{t+1}}} \; . \tag{13}$$

The proportionality absorbs a positive constant that depends only on the period t state of the system and is not affected by uncertainty or changes in the preference specification. A one percent increase in the marginal abatement cost Λ' increases the abatement rate μ by approximately half a percent. Under certainty, equation (13) states that the optimal abatement rate increases in the marginal value of climate capital ($-M$) and decreases in the marginal value of produced capital. Note that the prudence term cancels in equation (13), as it equally affects the marginal value of produced capital and climate capital.

The denominator on the right side (13) measures the pessimism weighted marginal value of capital under uncertainty. We found in section 4.4.1 that it always increases under uncertainty (for our Epstein-Zin-Weil scenarios). Here, it reduces the optimal abatement rate by increasing the opportunity value of investing in produced capital. From Figures 4-6, however, we know that optimal abatement increases for high values of η and decreases for low values of η. Both of these effects are mostly driven by the numerator in equation (13) and would also prevail in the absence of the produced capital uncertainty effect.

The marginal value of climate capital $-\frac{\partial V_{t+1}}{\partial M_{t+1}}$ is always positive. However, it is not a priori obvious how the marginal value of climate capital depends on growth shocks. Figure 12 in Appendix A shows that for $\eta = 2$ this marginal value of climate capital decreases in the technology level, while for $\eta = \frac{2}{3}$ it increases in the technology level. The figure also shows that this finding is independent of using Epstein-Zin-Weil preferences and even of the presence of growth uncertainty. The explanation of this asymmetry derives straight from our discussion in section 4.2 of how deterministic technological progress affects optimal abatement. Equation (9) expresses that abatement increases if (and only if) damages are more sensitive to production shocks than the marginal value of consumption: $\text{AIS}(u, c) < \text{Dam}_1(d, y)$. Translated into the value function, if this inequality is satisfied, then the marginal value of an abated ton of carbon (climate capital) increases in the technology level. That is the case for $\eta = \frac{2}{3} < 1$, but not for $\eta = 2 > 1$.

Epstein-Zin-Weil preferences give rise to the pessimism term that increases the weight on the low technology realizations. In the $\eta = 2$ scenario, low technology realizations imply a high marginal value of climate capital, increasing the expected marginal value of a carbon reduction. This increase of the marginal value of climate capital under uncertainty is over three times as large as the increase of the marginal value of produced capital. As a consequence, the value increase of climate capital (numerator in equation 13) dominates the opportunity cost of capital (denominator in equation 13), and the optimal abatement rate increases. In the $\eta = \frac{2}{3}$ scenario, the bad states of the world corresponding to low technology realizations also imply a lower marginal value of the climate capital. The pessimism term therefore reduces the expected value of climate capital, and at the same time increases the expected value of produced capital. Thus, pessimism weighting reduces the optimal abatement rate in the $\eta = \frac{2}{3}$ scenario (the numerator in equation (13) falls and the denominator increases), where damages are more sensitive to production shocks than the marginal value of consumption.

We close this section by explaining the abatement effect directly using our formula for the social cost of carbon in equation (7). We obtain the relevant intuition from interacting the pessimism term P_t with the two components that we already analyzed in the case of entangled preferences: marginal utility $u'(c)$ and damages $d(y)$. For qualitative understanding, we can restrict attention to the interactions $P_t(V(a))u'(c(a))d(y(a))$ within one period. By Jensen's inequality, these terms increase the social cost of carbon if and only if their product is convex. We focus on identifying and explaining the dominant effect in the resulting convexity condition

$$\frac{d^2}{da^2} P_t(V(a))\ u'(c(a))\ d(y(a)) = P_t(V(a)) \underbrace{\frac{d^2}{da^2}\Big(u'(c(a))d(y(a))\Big)}_{\text{term 1}} \tag{14}$$

$$+\frac{d}{da} P_t(V(a)) \underbrace{\frac{d}{da}\Big(u'(c(a))d(y(a))\Big)}_{\text{term 2}} + \frac{d^2}{da^2} P_t(V(a)) \underbrace{\Big(u'(c(a))d(y(a))\Big)}_{\text{term 3}} > 0\ .$$

Term 1 is multiplied by the positive constant[16] $P_t(V(a))$ and coincides otherwise with the term analyzed already for entangled preferences and characterized in equation (11). Term 3 is positive and multiplied with the positive constant $\frac{d^2}{da^2} P_t(V(a))$ arising from the convexity of the pessimism term. This third contribution always contributes positively to the convexity in equation (14) and, thus, always increases the optimal abatement rate under uncertainty. The dominating contribution is the pessimism weighted second term. The derivative $\frac{d}{da} P_t$ is negative, which is precisely the reason why it acts as a pessimism term: low realizations obtain a high weight. Term 2 itself is the same derivative that characterizes whether a deterministic increase in the technology parameter increases abatement (see section 4.2). The term is proportional to the condition discussed in equation (9): it is positive if, and only if, climate damages

[16] Note that we evaluate the convexity condition at its expected value. Earlier in this section we analyzed the pessimism term's fluctuations for different states of the world. Now we use Jensen's inequality and the analytic convexity characterization instead, allowing us to directly evaluate values, derivatives, and curvatures at a given point.

are more sensitive to growth uncertainty than marginal valuation. However, term 2 is weighted with the negative derivative of the pessimism term. Thus, the condition observed in equation (9) for deterministic growth flips its sign under uncertainty and Epstein-Zin-Weil preferences: Uncertainty increases optimal abatement if and only if $\text{AIS}(u, c) > \text{Dam}_1(d, y)$. In summary, what matters (most) under Epstein-Zin-Weil preferences is the pessimism-weighted expectation, which gives more weights to bad (welfare) states of the world. These bad states of the world are those corresponding to low technology realizations and, thus, combinations of relatively higher marginal utility and relatively lower physical damages. These states contribute most to increasing abatement under a strong aversion to intertemporal substitution (sharper increase of marginal utility) and when damages have a low sensitivity (modest decrease of physical damages).

5 Conclusions

Growth drives emissions and climate damages as well as well-being and the marginal valuation of damages. We analyze the consequences of growth uncertainty for optimal mitigation policy in a numerical model based on the integrated assessment model DICE. In an analytic discussion, we identify the relevant structural assumptions that drive the results. In the standard DICE model, with its usual preference specification, deterministic growth decreases the optimal present-day carbon tax. This result holds as long as the propensity to smooth consumption over time dominates the damage sensitivity to production. For uncertainty, the mechanism driving climate policy differ between the standard expected utility model and the more comprehensive risk pricing model that better accounts for risk premia and discount rates. Persistence in the growth shocks always increases the magnitude of the uncertainty effect.

Using a stochastic dynamic programming version of the DICE model with Nordhaus's (2008) preference calibration, we find that growth uncertainty slightly increases the optimal carbon tax. The effect is small, but contrasts with an earlier Monte-Carlo based study that suggests growth uncertainty reduces optimal abatement in DICE. The driving assumption is not the consumption smoothing parameter, which coincides with the Arrow-Pratt measure of relative risk aversion, but prudence, which captures the decrease of (absolute) risk aversion in income. Investment in produced capital always increases under growth uncertainty when prudence is positive, as is the case for isoelastic preferences used in the DICE model. In contrast, investment in climate capital increases under growth uncertainty only if prudence dominates the sensitivity of climate damages to production. If production is concave in technological progress, a high damage sensitivity weakens the requirements on prudence, and a high consumption smoothing preferences increases the prudence level where uncertainty implies emission reductions.

Our extension to comprehensive risk preferences implies much higher risk adjustments of the optimal carbon tax. Increasing the coefficient of relative risk aversion to its disentangled estimate in the finance literature implies an increase of the present optimal carbon tax by over 20% under an iid shock (to about \$40 per ton of carbon), and by over 45% under a moderately persistent shock (to about \$50). These uncertainty adjustments of the optimal carbon tax are several times larger than the risk premia found for damage uncertainty in a similar stochastic DICE model by Crost & Traeger (2010), and they are larger than the tax adjustments induced by Lemoine & Traeger's (2013f) extension to carbon cycle and feedback tipping points. However, the empirical asset pricing literature also suggests a reduction of the consumption smoothing parameter, explaining the low risk-free discount rate. Lowering the aversion to intertemporal fluctuations flips the sign of the uncertainty effect, causing uncertainty to decrease optimal abatement. The optimal carbon tax falls by 15% under iid shocks (to slightly above \$70) and by 30% under moderate shock persistence (to below \$60). The overall carbon tax is still higher in this scenario because the deterministic increase of the optimal carbon tax resulting from the lower consumption smoothing parameter is greater than the uncertainty effect.

An endogenous pessimism weighting drives the large uncertainty effects in the comprehensive preference framework. The pessimism term arises from risk aversion dominating consumption smoothing, and biases upwards the weights on states of the world that lead to lower welfare. In our context, these states are the low growth states. Low growth states always correspond to a higher productivity of produced capital and, as a consequence, the investment rate in produced capital increases under uncertainty by approximately 1 percentage point (2-5%) in all disentangled preference scenarios. Low growth states increase the marginal value of climate capital only if consumption smoothing dominates damage sensitivity, a condition satisfied for a high aversion to intertemporal substitution. If aversion to intertemporal substitution is low, as estimated in disentangled models, then damage sensitivity dominates; the pessimism bias increases the weight on scenarios with a relatively low value of climate damages, and uncertainty decreases the optimal abatement rate. The sign of the dominating pessimism contribution in the disentangled scenario depends, similarly to the deterministic case, on consumption smoothing and damage sensitivity to production. However, the pessimism bias implies that the uncertainty effect has the opposite sign on the optimal carbon tax as compared to a deterministic increase in the technology level.

A robust policy message is that all of our empirical simulations give rise to a higher optimal carbon tax under growth uncertainty. Quantitatively, this increase is significant when disentangling risk aversion and risk premia form intertemporal consumption smoothing and the risk-free discount rate. A second message is that under persistent growth shocks the optimal carbon tax becomes less sensitive to the consumption smoothing parameter. These results hold in the DICE model, which is was used in the US social cost of carbon assessment. Conceptually, we find that the policy impact of deterministic growth and growth uncertainty changes qualitatively in the neighborhood where "η" is close to unity, however, the interpretation differs. In a deterministic setting, it means that intertemporal consumption smoothing dominates the damage elasticity to production. Under uncertainty it means that prudence dominates twice the damage elasticity to production. Moreover, the uncertainty condition changes when damages are not linear in production or technological progress has a non-linear impact on production. Also the mechanism driving sign changes under uncertainty differs between the economic standard model and the disentangled asset pricing based model. Our numeric results employs fully rational, observed preferences. In the context of climate change, future wealth is the wealth consumed by future generations not currently alive. Instead of employing observed preferences, several scholars argue for the use of ethical evaluation criteria. Then, η in its interpretation as intertemporal, i.e., intergenerational consumption smoothing would be much larger. Less immediate is how ethical arguments change prudence and risk aversion, which are conceptually independent, and we leave this question for future research.

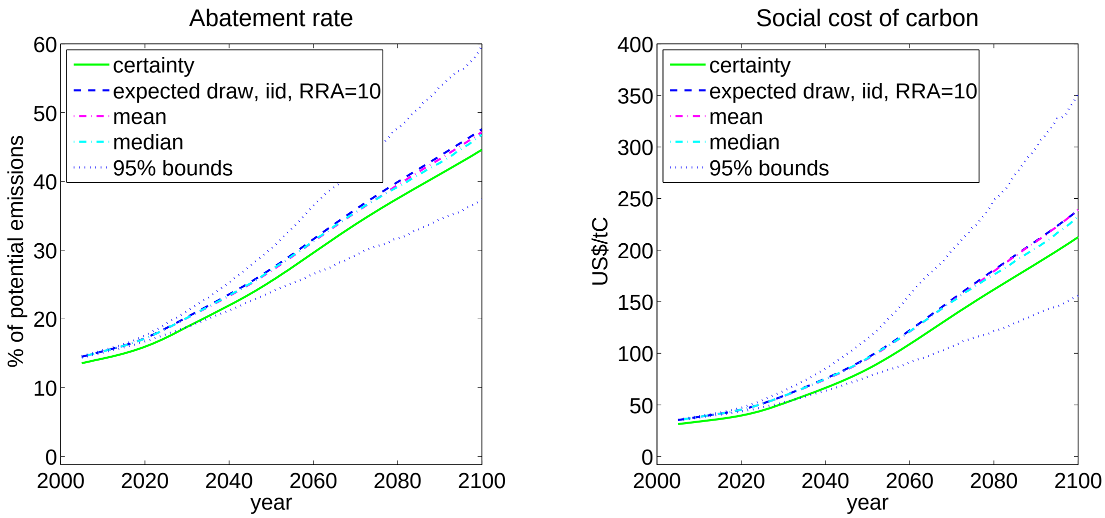

Figure 7 shows the mean, the median, the expected draw representation, and the 95 % confidence bounds of 1000 optimal random paths for abatement and the social cost of carbon. The social preference parameters are $\eta = 2, RRA = 10$. We compare the different measures to the the optimal paths under certainty. We observe that mean, meadian, and expected draw representation mostly coincide. The confidence intervals reveal considerable variation in the optimal climate policy in response to resolving growth uncertainty.

Appendix

A Further Results

Our figures in the main text represent uncertainty by expected draws. Figure 7 shows that these expected draw representations closely resemble the mean and the median policy of 1000 random path realizations. The uncertainty in the optimal policies is sizable. At the end of the century, there is a 5% chance that the abatement rate is lower than 38% or higher than 65% (with a median of 48%). The social cost of carbon lies with 95% confidence between \$160 and \$400.

Deterministic growth rate variations strongly influence optimal climate policy under certainty. Figure 8 shows three deterministic growth rates (for $\eta = 2$): The original DICE-2007 value, a growth rate that is 0.5 percentage points lower each year than the DICE value, and a growth rate that is 0.5 percentage points higher each year. The left panel in Figure 8 shows the optimal abatement rate and the right panel shows the optimal social cost of carbon. The higher the deterministic growth rate, the lower the optimal present day carbon tax: The marginal value of consumption decreases faster than the marginal damage of additional emissions (opposite of inequality (9)). But high productivity growth enables fast capital and wealth accumulation, increasing the relative valuation of emissions. Therefore the abatement rate in the high growth scenario eventually increases steeply and surpasses the abatement rates from the other scenarios. At the end of the century the high-growth abatement rate is more than 6 percentage points higher compared to the DICE-2007 baseline. In the low growth scenario, the abatement rate reaches 100% more than 100 years

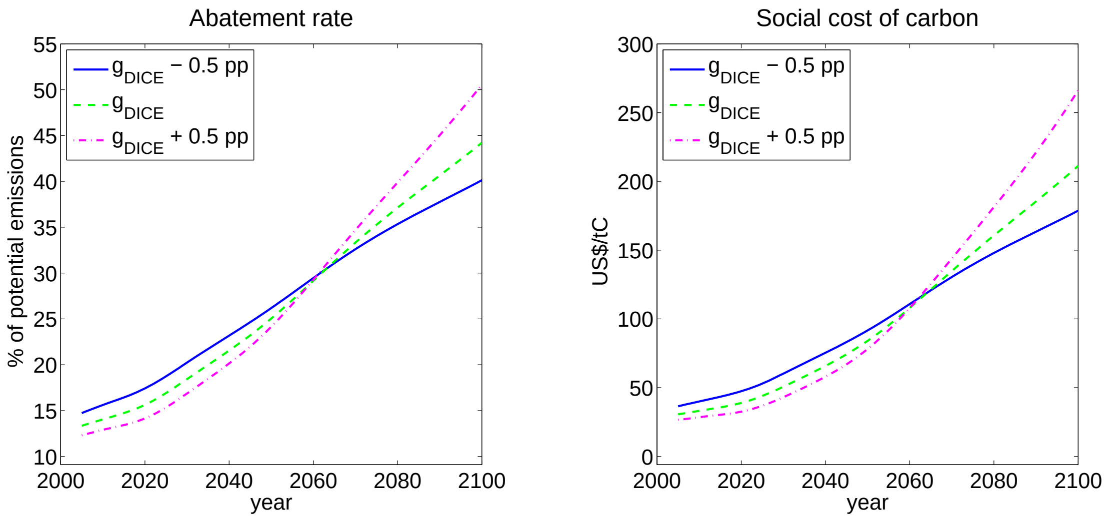

Figure 8 compares the optimal abatement rate and social cost of carbon under certainty for three growth rates: the DICE growth rate, the DICE growth rate +0.5 percentage points, and the DICE growth rate −0.5 percentage points. Higher growth lowers the optimal present-day abatement but increases abatement steeply later in the century.

later than in the high growth scenario (not shown).

Figure 9 shows the optimal consumption level and the corresponding consumption rate over the relative deviation of the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. Technological growth is uncertain and the decision-maker has standard entangled preferences (η = RRA = 2). The figure depicts 20 different lines corresponding to 20 different points in time along the optimal path. Time, carbon stock and capital are hence held constant, while normalized technology varies over its en-

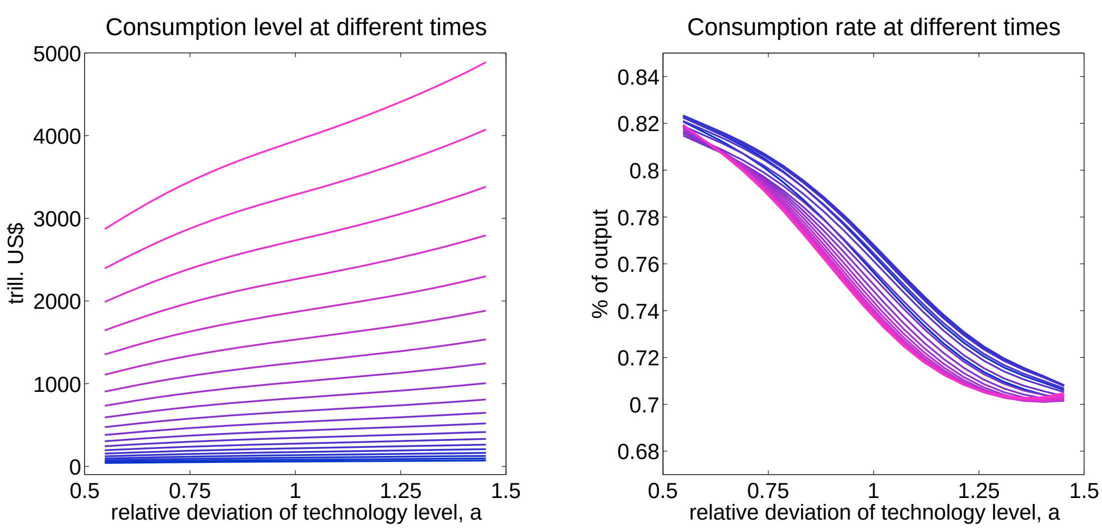

Figure 9 shows optimal consumption and the consumption rate over the normalized technology level ($a = \frac{A}{A^{det}}$) under uncertainty with entangled preferences (η = RRA = 2). The 20 different lines correspond to different points in time along the optimal path. Time intervals are evenly spaced from 2005 (dark blue) to 2400 (light pink). While the consumption level increases slightly in technology level, the consumption rate decreases.

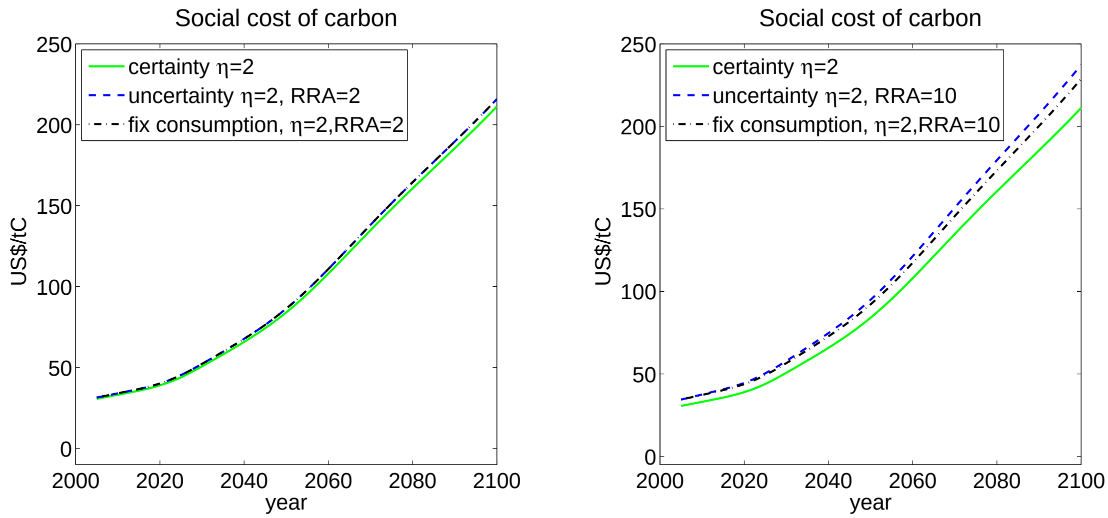

Figure 10 shows the optimal carbon tax under certainty, uncertainty, and uncertainty with the consumption rate fixed at the deterministically optimal level. The left panel displays standard entangled preferences (η = RRA = 2), the right panel shows the disentangled preference scenario ($\eta = 2$ and RRA = 10). For standard preferences, fixing consumption to its deterministic level has no notable impact on abatement; for disentangled preferences we observe a slightly lower social cost of carbon.

tire range. The time intervals are evenly spaced between 2005 (dark blue) and 2400 (light pink). While the consumption level increases slightly in the technology level, the consumption rate decreases.

Figure 10 displays the optimal carbon tax under certainty, uncertainty, and for a case in which we fix the consumption rate at the deterministically optimal level (only abatement responds to uncertainty). The left panel displays the case of standard entangled preferences (η = RRA = 2), and the right panel shows the disentangled preferences scenario ($\eta = 2$ and RRA = 10). For standard preferences, fixing consumption to its deterministic level has no notable impact on abatement. For disentangled preferences we observe a slightly lower social cost of carbon.

In Figure 11 we plot the normalized value function over the relative deviation of the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. We show the two disentangled preference scenarios ($\eta = 2, RRA = 10$ and $\eta = 2/3, RRA = 10$). As in Figure 9, the figure depicts 20 lines, each of which corresponds to a different point in time. We observe that the value function is significantly more concave in the case $\eta = 2$ than in the case $\eta = 2/3$.

In Figure 12 we show the marginal welfare gain $-\frac{\partial V_t^*}{\partial M_t}$ from an avoided ton of atmospheric carbon as a function of the relative deviation of the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. If the preference for consumption smoothing is strong ($\eta = 2$, left panels) the marginal gain decreases in the technology level, if the desire to smooth consumption is weak ($\eta = 2/3$, right panels) the marginal welfare gain increases in the technology level. This finding holds under certainty (upper panels) as well as uncertainty (lower panels). The figure again depicts 20 lines, each of which corresponds to a different point in time.

Figures 13 and 14 show the result of calibrating our simplified climate module

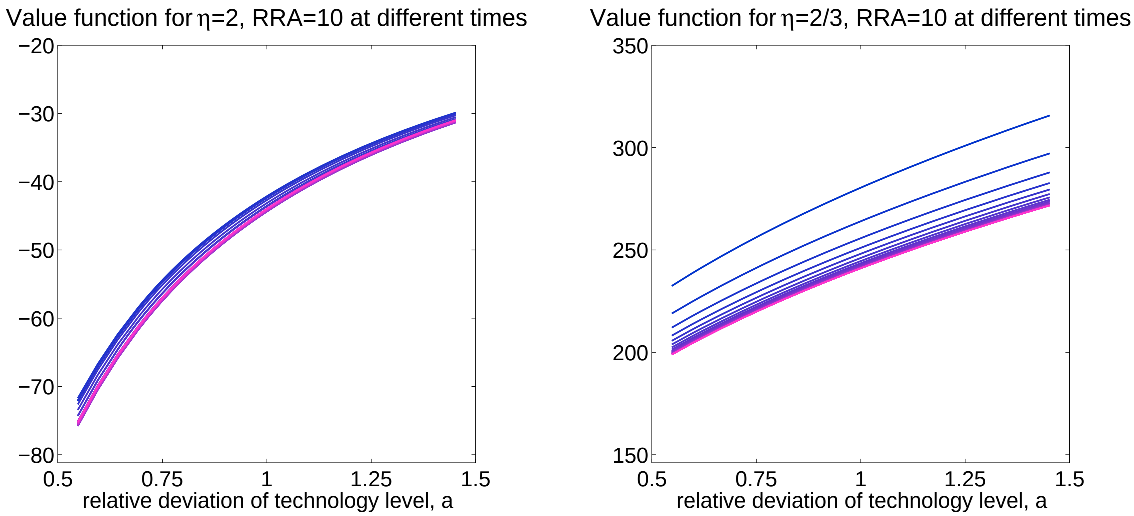

Figure 11 shows the normalized value function over the relative deviation of the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. We show two levels of consumption smoothing ($\eta = 2$, left panel and $\eta = 2/3$, right panel, both with RRA = 10). The 20 different lines correspond to different points in time along the optimal path. Time intervals are evenly spaced from 2005 (dark blue) to 2400 (light pink). We observe that the value function is significantly more concave in the case $\eta = 2$ than in the case $\eta = 2/3$.

to the original DICE-2007 model. Figure 13 shows the case of standard preferences ($\eta = 2$), whereas in Figure 14 the desire to smooth consumption is relatively low ($\eta = 2/3$). The calibration is the same for both sets of graphs and the differences are similar. The optimal climate policies (abatement rate and carbon tax) and the evolution of the carbon stock resemble DICE closely. To calibrate these well, we accept a slightly larger deviation of temperature.

B Renormalizing 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. 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:[17]

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

[17]For the sake of clarity, some equations from section 2.3 are reproduced here.

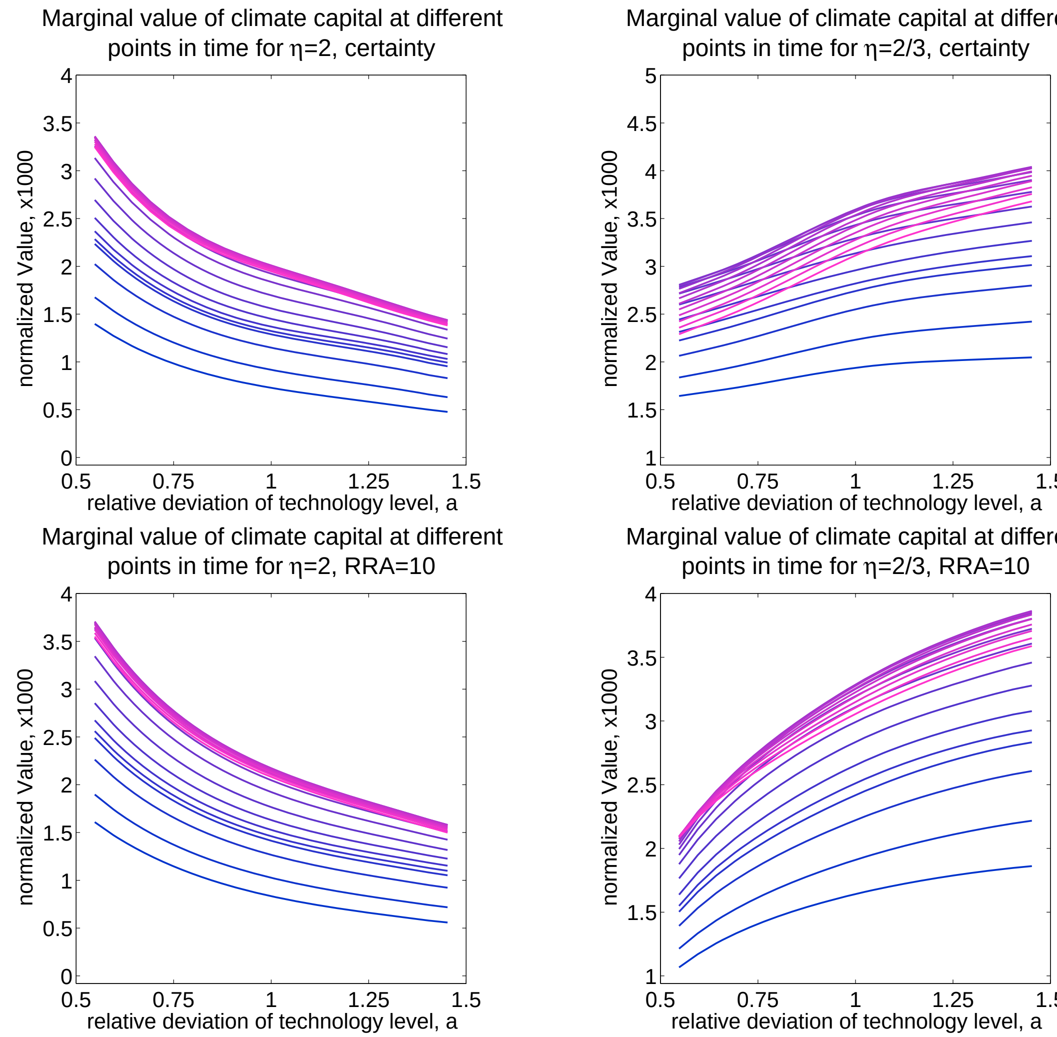

Figure 12 shows the marginal welfare gain $-\frac{\partial V_t^*}{\partial M_t}$ from an avoided ton of atmospheric carbon as a function of the normalized technology level $a = \frac{A}{A^{det}}$. In the case of a relatively high preference for consumption smoothing ($\eta = 2$, left panels) the marginal welfare gain decreases in the technology level. In the case of a relatively low preference for consumption smoothing ($\eta = 2/3$, right panels) the marginal welfare gain increases in the technology level. This finding holds under certainty (upper panels) as well as uncertainty (lower panels). The figure depicts 20 lines, each of which corresponds to a different point in time. Time intervals are evenly spaced from 2005 (dark blue) to 2400 (light pink).

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. We normalize by the deterministic technology level A^{det}, the level of technology in the certainty scenario (with all shocks equal zero, $z_t = 0 \ \forall t$)

$$A_{t+1}^{det} = A_t^{det} \exp\left[\bar{g}_{A,t}\right] \quad \text{where} \quad \bar{g}_{A,t} = g_{A,0} \exp\left[\delta_A \cdot t\right] \ .$$

Expressing consumption and capital in effective labor terms results in the definitions $c_t = \frac{C_t}{A_t^{det} L_t}$ and $k_t = \frac{K_t}{A_t^{det} L_t}$. Moreover, we define $a_t = \frac{A_t}{A_t^{det}}$. The normalized productivity one period ahead is then defined as

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

With the normalized variables we transform the Bellman equation (5):

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

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

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

which leads to the new Bellman equation (6)

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

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 18). This switch of variables turns the constraints on the optimization problem linear.

We recover the original value function from

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

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

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

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

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

C Derivation of Analytic Formulas

Here we derive the main equations of section 4. Let α denote the constant consumption rate.

Derivation of equation (9):
The optimal carbon tax increases under deterministic technological progress if $\frac{d}{da}u'(c(a))d(y(a))$ is positive

$$\begin{aligned}
\frac{d}{da}u'(\alpha y(a))d(y(a)) &= u''(\alpha y(a))\alpha y'(a)d(y(a)) + u'(\alpha y(a))d(y(a))y'(a) \\
&= \left[\frac{u''(\alpha y(a))}{u'(\alpha y(a))}\alpha y(a) + \frac{d'(y(a))}{d(y(a))}y(a)\right] u'(\alpha y(a))d(y(a))\frac{y'(a)}{y(a)} \\
&\propto \left[-\,\mathrm{MU}_1(u,c) + \mathrm{Dam}_1(d,y)\right],
\end{aligned}$$

given that marginal utility, damages, and the relation between production and technology level are positive.

Derivation of equation (10):
By Jensen's inequality, uncertainty over production increases abatement if the product $u'(\alpha y)d(y)$ is convex:

$$\frac{d^2}{dy^2}u'(\alpha y)d(y) = \alpha^2 u'''(y)d(y) + 2\alpha u''(y)d'(y) + u'(y)d''(y) > 0$$

Assuming a positive propensity to smooth consumption over time ($u''(\alpha y) < 0$) and positivity of damages, we can rewrite the condition as

$$-\frac{u'''(\alpha y)}{u''(\alpha y)}\alpha y = -\frac{u'''(c)}{u''(c)}c > 2\frac{d'(y)}{d(y)}y - \frac{\frac{d''(y)}{d(y)}y^2}{-\frac{u''(\alpha y)}{u'(\alpha y)}\alpha y} = 2\frac{d'(y)}{d(y)}y - \frac{\frac{d''(y)}{d'(y)}y\frac{d'(y)}{d(y)}y}{-\frac{u''(c)}{u'(c)}c} ,$$

which coincides with equation (10).

Derivation of equation (11):
By Jensen's inequality, growth uncertainty increases abatement if the product $u'(\alpha y(a))d(y(a))$

is convex:

$$\frac{d^2}{dy^2}u'(\alpha y)\, d(y) = \alpha^2 u'''(\alpha y)\left[y'(a)^2\right] d(y)$$
$$+\alpha u''(\alpha y)(\alpha y)y''(a)d'(y) + 2\alpha u''(\alpha y)(\alpha y)\left[y'(a)^2\right] d'(y)$$
$$+u'(\alpha y)d''(y)\left[y'(a)^2\right] + u'(\alpha y)d''(y)y''(a) > 0$$

Assume a positive propensity to smooth consumption and positivity of damages. Using $c = \alpha y(a)$, we can rewrite this condition as

$$-\frac{u'''(c)}{u''(c)}cy'(a)^2 - y''(a)y - 2y'(a)^2\frac{d'(y)}{d(y)}y + \frac{\frac{d''(y)}{d(y)}y^2}{-\frac{u''(c)}{u'(c)}c}y'(a)^2 - \frac{u'(c)y^2}{u''(c)c}\frac{d'(y)}{d(y)}y''(a) > 0$$

$$\Rightarrow \text{Prud}(u,c)\; y'(a)^2 > 2\,\frac{d'(y)}{d(y)}y\; y'(a)^2 - \frac{\frac{d''(y)}{d(y)}y^2}{\text{AIS}(u,c)}y'(a)^2 + y''(a)y - y''(a)y\frac{\frac{d'(y)}{d(y)}y}{\text{AIS}(u,c)}$$

$$\Rightarrow \text{Prud}(u,c) > 2\,\frac{d'(y)}{d(y)}y - \frac{\frac{d''(y)}{d(y)}y^2}{\text{AIS}(u,c)} + \frac{y''(a)y}{y'(a)^2} - \frac{y''(a)y}{y'(a)^2}\frac{\frac{d'(y)}{d(y)}y}{\text{AIS}(u,c)}$$
$$> \text{Dam}_1(d,y)\left[2 - \frac{\text{Dam}_2(d,y)}{\text{AIS}(u,c)}\right] + \frac{\text{Tech}_2(\text{y},a)}{\text{Tech}_1(\text{y},a)}\left[1 - \frac{\text{Dam}_1(d,y)}{\text{AIS}(u,c)}\right]$$

The last equation coincides with equation (11) in section 4.3.2.

Derivation of equation (12):
Optimizing the normalized Bellman equation (6) with respect to consumption returns

$$u'(c_t) = \beta_t \underbrace{\exp(-g_{A,t} - g_{L,t})}_{\equiv g_t} \underbrace{\frac{\mathbf{E}_t f'(V_{t+1})}{f'(f^{-1}\mathbf{E}_t f(V_{t+1}))}}_{\equiv \Pi_t}\; \mathbf{E}_t \underbrace{\frac{f'(V_{t+1})}{\mathbf{E}_t f'(V_{t+1})}}_{\equiv P_t} \frac{\partial V_{t+1}}{\partial k_{t+1}} \qquad (15)$$

In (12) we use the definitions for pessimism (P_t) and prudence (Π_t) and the proportionality absorbs the growth factor g_t.

Derivation of equations (8) **and** (13):

Optimizing the normalized Bellman equation (6) with respect to abatement returns

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

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

$$\Rightarrow \Lambda'(\mu_t) = -\alpha(T_t, t) \ \beta_t \Pi_t \frac{\mathbf{E}_t \ P_t \ \frac{\partial V_{t+1}}{\partial M_{t+1}}}{u'(c_t)} \ , \tag{16}$$

where we use equation (15) in the last step. The second line corresponds to equation (13) in which the proportionality disguises the first term on the right hand side. Differentiating the Bellman equation (6) partially with respect to the carbon stock M_t using the envelope theorem returns

$$\frac{\partial V_t}{\partial M_t} = \beta_t \Pi_t \mathbf{E}_t P_t \left[\frac{\partial V_{t+1}}{\partial M_{t+1}} \underbrace{\left[(1-\delta_{M,t}) + \frac{\partial \delta_{M,t}}{\partial M_t}(M_t - M_{pre}) \right]}_{\frac{\partial M_{t+1}}{\partial M_t}} + \frac{\partial V_{t+1}}{\partial k_{t+1}} g_t \frac{\partial y_t}{\partial M_t} \right]$$

$$= u'(c_t) \frac{\partial y_t}{\partial M_t} + \beta_t \frac{\partial M_{t+1}}{\partial M_t} \Pi_t \mathbf{E}_t P_t \frac{\partial V_{t+1}}{\partial M_{t+1}} \ ,$$

where we again use equation (15). Repeated substitution of this relation advancing the time indices by one period implies

$$\frac{\partial V_t}{\partial M_t} = u'(c_t) \frac{\partial y_t}{\partial M_t} + \beta_t \frac{\partial M_{t+1}}{\partial M_t} \Pi_t \mathbf{E}_t P_t \ \ u'(c_{t+1}) \frac{\partial y_{t+1}}{\partial M_{t+1}} +$$

$$\beta_t \frac{\partial M_{t+1}}{\partial M_t} \Pi_t \mathbf{E}_t P_t \quad \beta_{t+1} \frac{\partial M_{t+2}}{\partial M_{t+1}} \Pi_{t+1} \mathbf{E}_{t+1} P_{t+1} \ \ \frac{\partial V_{t+2}}{\partial M_{t+2}}$$

$$= u'(c_t) \frac{\partial y_t}{\partial M_t} + \sum_{\tau=t}^{\infty} \left\{ \prod_{j=t}^{\tau} \beta_j \frac{\partial M_{j+1}}{\partial M_j} \Pi_j \mathbf{E}_j P_j \right\} \ \ u'(c_{\tau+1}) \frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \ . \tag{17}$$

Inserting equation (17) into equation (16) gives us

$$\Lambda'(\mu_t) = -\alpha(T_t,t)\Bigg[\frac{\beta_t \Pi_t \mathbf{E}_t \; P_t \; u'(c_{t+1})\frac{\partial y_{t+1}}{\partial M_{t+1}}}{u'(c_t)} +$$

$$\frac{\beta_t \Pi_t \mathbf{E}_t \; P_t \sum_{\tau=t+1}^{\infty}\left\{\prod_{j=t+1}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j \mathbf{E}_j P_j\right\} \;\; u'(c_{\tau+2})\frac{\partial y_{\tau+2}}{\partial M_{\tau+2}}}{u'(c_t)}\Bigg]$$

$$= -\frac{\alpha(T_t,t)}{\frac{\partial M_{t+1}}{\partial M_t}}\sum_{\tau=t}^{\infty}\left\{\prod_{j=t}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j \mathbf{E}_j P_j\right\} \;\; \frac{u'(c_{\tau+1})}{u'(c_t)}\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}$$

$$= -\frac{\alpha(T_t,t)}{\frac{\partial M_{t+1}}{\partial M_t}}\,\mathbf{E}_t^* \,\sum_{\tau=t}^{\infty}\left\{\prod_{j=t}^{\tau}\beta_j \frac{\partial M_{j+1}}{\partial M_j}\Pi_j P_j\right\} \;\; \frac{u'(c_{\tau+1})}{u'(c_t)}\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}} \;.$$

In equation (8) we multiplied out $\frac{\partial M_\tau}{\partial E_t} = \prod_{j=t}^{\tau}\frac{\partial M_{j+1}}{\partial M_j}$ and the proportionality absorbs the first term on the right hand side. While the expectation operators $\mathbf{E}_t$ take expectations over the realization of $\tilde{A}_{t+1}$ (precisely the normalized $\tilde{a}_{t+1}$) conditional on earlier realizations of A_t, the operator $\mathbf{E}_t^*$ takes expectations over all possible future sequences $\tilde{A}_{t+1}, \tilde{A}_{t+2}, \ldots$ conditional on A_t.

D The Climate Enriched Economy Model

The following model is largely a reproduction of DICE-2007 (Nordhaus 2008). 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. For a detailed discussion of the changes to DICE refer to Traeger (2012). His model differs only by including an extra temperature state to account temperature delay. All parameters are characterized and quantified in Table D on page 47.

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 mimic 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 13 and 14).

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 (μ_t). 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. We discuss the uncertain technological progress, given by equation (2) in detail 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

$$\Lambda_t(\mu_t) = \Psi_t \mu_t^{a_2} \tag{18}$$

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 13 and 14 compare the performance of our model to the original DICE model.

E Growth Rate Shocks

Suppose that t is the period characterizing the information of the decision-maker, i.e. she takes expectation over the future as if being in period t. Recall that we denote by

A_t the current technology level at time t and by $\tilde{A}_{t+1}$ the uncertain future technology level. Let A^{det}_{t+1} be the hypothetical deterministic technology level; the technology level without shocks in any period. Finally, $a_t = A_t/A^{det}_t$ is the normalized technology level, the actual level as a multiple of the deterministic level. Technological growth has a deterministic and a random component, so that the technology level one period ahead is:

$$\tilde{A}_{t+1} = A_t \exp\left[\tilde{g}_{A,t}\right] = A_t \exp[\bar{g}_{A,t} + \tilde{z}_t] \ ,$$

where $\bar{g}_{A,t}$ is the deterministic growth trend and the growth shock $\tilde{z}$ is specified further below. Then the normalized technology level one period ahead is:

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

From the perspective of period t, the technology level in $t+\tau$ is cumulative in the growth shocks:

$$\begin{aligned} \tilde{A}_{t+\tau} &= \tilde{a}_{t+\tau} A^{det}_{t+\tau} = \exp[\tilde{z}_{t+\tau-1}] \tilde{a}_{t+\tau-1} A^{det}_{t+\tau} \\ &= \exp[\tilde{z}_{t+\tau-1} + \tilde{z}_{t+\tau-2}] \tilde{a}_{t+\tau-2} A^{det}_{t+\tau} \\ &= \exp\left[\textstyle\sum_{\tau'=0}^{\tau-1} \tilde{z}_{t+\tau'} \right] a_t A^{det}_{t+\tau} \ . \end{aligned} \tag{19}$$

E.1 Iid Shocks

First we consider a growth shock that is normally, iid distributed

$$\tilde{z}_t = \tilde{x}_t \sim \mathcal{N}(\mu_x, \sigma_x^2) \ .$$

Technololgy is hence lognormally distributed. Taking expectations in equation (19) for this shock gives

$$\begin{aligned} \mathbf{E}\tilde{A}_{t+\tau} &= \mathbf{E} \exp\left[\textstyle\sum_{\tau'=0}^{\tau-1} \tilde{x}_{t+\tau'} \right] a_t A^{det}_{t+\tau} \\ &= \exp\left[\textstyle\sum_{\tau'=0}^{\tau-1} \mu_A + \frac{\sigma_A^2}{2} \right] a_t A^{det}_{t+\tau} \\ &= \exp\left[\tau \left(\mu_A + \frac{\sigma_A^2}{2} \right) \right] a_t A^{det}_{t+\tau} \ . \end{aligned}$$

Thus, setting $\mu_x = -\frac{\sigma_x^2}{2}$ implies $\exp[\tau(\mu_A + \frac{\sigma_A^2}{2})] = 1$ and equates the $\tilde{A}_{t+\tau}$ expectations with the hypothetical development under certainty from t onwards. Note that the 'cumulative' variance of the normal distribution in the exponent for $\tilde{A}_{t+\tau}$ increases linearly over time.

E.2 Persistent Shocks

Now consider shocks that affect the economy for more than one period. Set

$$\tilde{z}_t = \tilde{x}_t + \tilde{y}_t \quad \text{where} \quad \tilde{y}_t = \gamma y_{t-1} + \tilde{\epsilon}_t \quad \text{and} \quad \tilde{\epsilon}_t \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2) \ ,$$

where $\tilde{x}_t$ is iid normally distributed again. To calculate the expectation, we need to pick a value for past shocks y_{t-1}. We assume $y_{t-1} = 0$. The random variable $\tilde{y}_{t+\tau}$ can be written as

$$\tilde{y}_{t+\tau} = \gamma^\tau y_t + \sum_{i=1}^{\tau} \gamma^{\tau-i} \tilde{\epsilon}_{t+i} \ . \tag{20}$$

Inserting $\tilde{x}_t + \tilde{y}_t$ in (19), the expectation for the technology level multiple periods ahead $\tilde{A}_{t+\tau}$ is

$$\begin{aligned}\mathbf{E}\tilde{A}_{t+\tau} &= \mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1} \tilde{x}_{t+\tau'} + \tilde{y}_{t+\tau'}\right] a_t A_{t+\tau}^{det} \\ &= \mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1} \tilde{x}_{t+\tau'}\right] \cdot \mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1} + \tilde{y}_{t+\tau'}\right] a_t A_{t+\tau}^{det} \ .\end{aligned}$$

$\mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1} x_{t+\tau'}\right] = 1$ for $\mu_x = -\frac{\sigma_x^2}{2}$ as shown in above in section E.1. Inserting from (20)

$$\begin{aligned}\mathbf{E}\tilde{A}_{t+\tau} &= \mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1}\left[\gamma^{\tau'} y_t + \sum_{i=0}^{\tau'-1} \gamma^{\tau'-i-1}\ \tilde{\epsilon}_{t+i}\right]\right] a_t A_{t+\tau}^{det} \\ &= \exp\left[\sum_{\tau'=0}^{\tau-1} \gamma^{\tau'} y_t\right] \mathbf{E}\exp\left[\sum_{\tau'=0}^{\tau-1} \sum_{i=0}^{\tau'-1} \gamma^{\tau'-i-1}\ \tilde{\epsilon}_{t+i}\right] a_t A_{t+\tau}^{det} \\ &= \exp\left[\sum_{\tau'=0}^{\tau-1} \gamma^{\tau'} y_t\right] \mathbf{E}\exp\left[\sum_{i=0}^{\tau}\left[\sum_{j=0}^{\tau-i-1} \gamma^j\right]\ \tilde{\epsilon}_{t+i}\right] a_t A_{t+\tau}^{det} \\ &= \exp\left[\sum_{\tau'=0}^{\tau-1} \gamma^{\tau'} y_t\right] \mathbf{E}\exp\left[\sum_{i=0}^{\tau} \frac{1-\gamma^{\tau-i}}{1-\gamma}\ \tilde{\epsilon}_{t+i}\right] a_t A_{t+\tau}^{det} \\ &= \exp\left[\sum_{\tau'=0}^{\tau-1} \gamma^{\tau'} y_t\right] \exp\left[\sum_{i=0}^{\tau} \frac{1-\gamma^{\tau-i}}{1-\gamma}\ \left(\mu_\epsilon + \frac{\sigma_\epsilon^2}{2}\right)\right] a_t A_{t+\tau}^{det} \ .\end{aligned} \tag{21}$$

Thus, conditional on $y_t = 0$ setting $\mu_\epsilon = -\frac{\sigma_\epsilon^2}{2}$ equates the $\tilde{A}_{t+\tau}$ expectation with the hypothetical value that would result from growing at the deterministic rate $\tilde{g}_{A,t+\tau}$ from t onwards.

Equation (21) tells us how the 'cumulative' variance of the normal distribution in the exponent for $\tilde{A}_{t+\tau}$ increases over time. The factors $^{1-\gamma^{\tau-i+1}}/_{1-\gamma} > 1$ increase in τ so that the "aggregate variance" increases more than linearly. Uncertainty over the next period capital stock conditional on $y_t = 0$ is the same as in the iid scenario, but looking further into the future uncertainty increases more in the case of persistence.

References

Ackerman, F., Stanton, E. A. & Bueno, R. (2010), 'Fat tails, exponents, extreme uncertainty: Simulating catastrophe in DICE', *Ecological Economics* **69**(8), 1657–1665.

Anthoff, D. & Tol, R. S. J. (2009), 'The impact of climate change on the balanced growth equivalent: An application of FUND', *Environmental and Resource Economics* **43**, 351 – 367.

Anthoff, D., Tol, R. S. J. & Yohe, G. W. (2009), 'Risk aversion, time preference, and the social cost of carbon', *Environmental Research Letters* **4**, 1–7.

Arouba, S. B., Fernandez-Villaverde, J. & Rubio-Ramirez, J. F. (2006), 'Comparing solution methods for dynamic equilibrium economies', *Journal of Economic Dynamics and Control* **30**, 2477–2508.

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**, 2817–2828.

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., Kiku, D. & Yaron, A. (2012), 'An empirical evaluation of the long-run risks model for asset prices', *Critical Finance Review* **1**, 183–221.

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.

Brock, W. A. & Mirman, L. J. (1972), 'Optimal economic growth and uncertainty: The discounted case', *Journal of Economic Theory* **4**, 479–513.

Chatterjee, S., Giuliano, P. & Turnovsky, S. J. (2004), 'Capital income taxes and growth in a stochastic economy: A numerical analysis of the role of risk aversion and intertemporal substitution', *Journal of Public Economic Theory* **6**(2), 277–310.

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

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.

Crost, B. & Traeger, C. P. (2013), 'Optimal climate policy: Uncertainty versus Monte-Carlo', *Economic Letters* **120**(3), 552–558.

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.

Epstein, L. G. & Zin, S. E. (1991), 'Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis', *Journal of Political Economy* **99**(2), 263–86.

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* (17348).

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* .

Jensen, S. & Traeger, C. (2013), 'Optimally climate sensitive policy: A comprehensive evaluation of uncertainty & learning', *Department of Agricultural and Resource Economics, UC Berkeley* .

Kaltenbrunner, G. & Lochstoer, L. (2010), 'Long-run risk through consumption smoothing', *Review of Financial Studies* **28**(8), 3190–3224.

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

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

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.

Kopp, R. E., Golub, A., Keohane, N. O. & Onda, C. (2012), 'The influence of the specification of climate change damages on the social cost of carbon', *Economic E-Journal* **6**, 1–40.

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

Lemoine, D. & Traeger, C. P. (2013f), 'Watch your step: Optimal policy in a tipping climate', *American Economic Journal: Economic Policy* (forthcoming).

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., Sergeyev, D. & Steinsson, J. (2012), 'Growth-rate and uncertainty shocks in consumption: Cross-country evidence', *NBER Working Paper* (18128).

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. (2011), 'Estimates of the social cost of carbon: Background and results from the RICE-2011 model', *Cowles Foundation Discussion Paper* (1826).

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

Pycroft, J., Vergano, L., Hope, C. W., Paci, D. & Ciscar, J. C. (2011), 'A tale of tails: Uncertainty and the social cost of carbon dioxide', *Economic E-Journal* **5**, 1–29.

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

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

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

Traeger, C. (2013), 'Why uncertainty matters - discounting under intertemporal risk aversion and ambiguity', *Economic Theory* (forthcoming).

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

Traeger, C. P. (2012), 'A 4-stated dice: Quantitatively addressing uncertainty effects in climate change', *CUDARE working paper 1130* .

Treasury, H. (2003), *The Green Book – Appraisal and Evaluation in Central Government*, HM Treasury, London.

Turnovsky, S. J. & Smith, W. T. (2006), 'Equilibrium consumption and precautionary savings in a stochastically growing economy', *Journal of Economic Dynamics and Control* **30**(2), 243–278.

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

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

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

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.

Weitzman, M. L. (2001), 'Gamma discounting', *The American Economic Review* **91**(1), pp. 260–271.

Weitzman, M. L. (2010), 'What is the damages function for global warming – and what difference might it make?', *Climate Change Economics* **1**(1), 57–69.

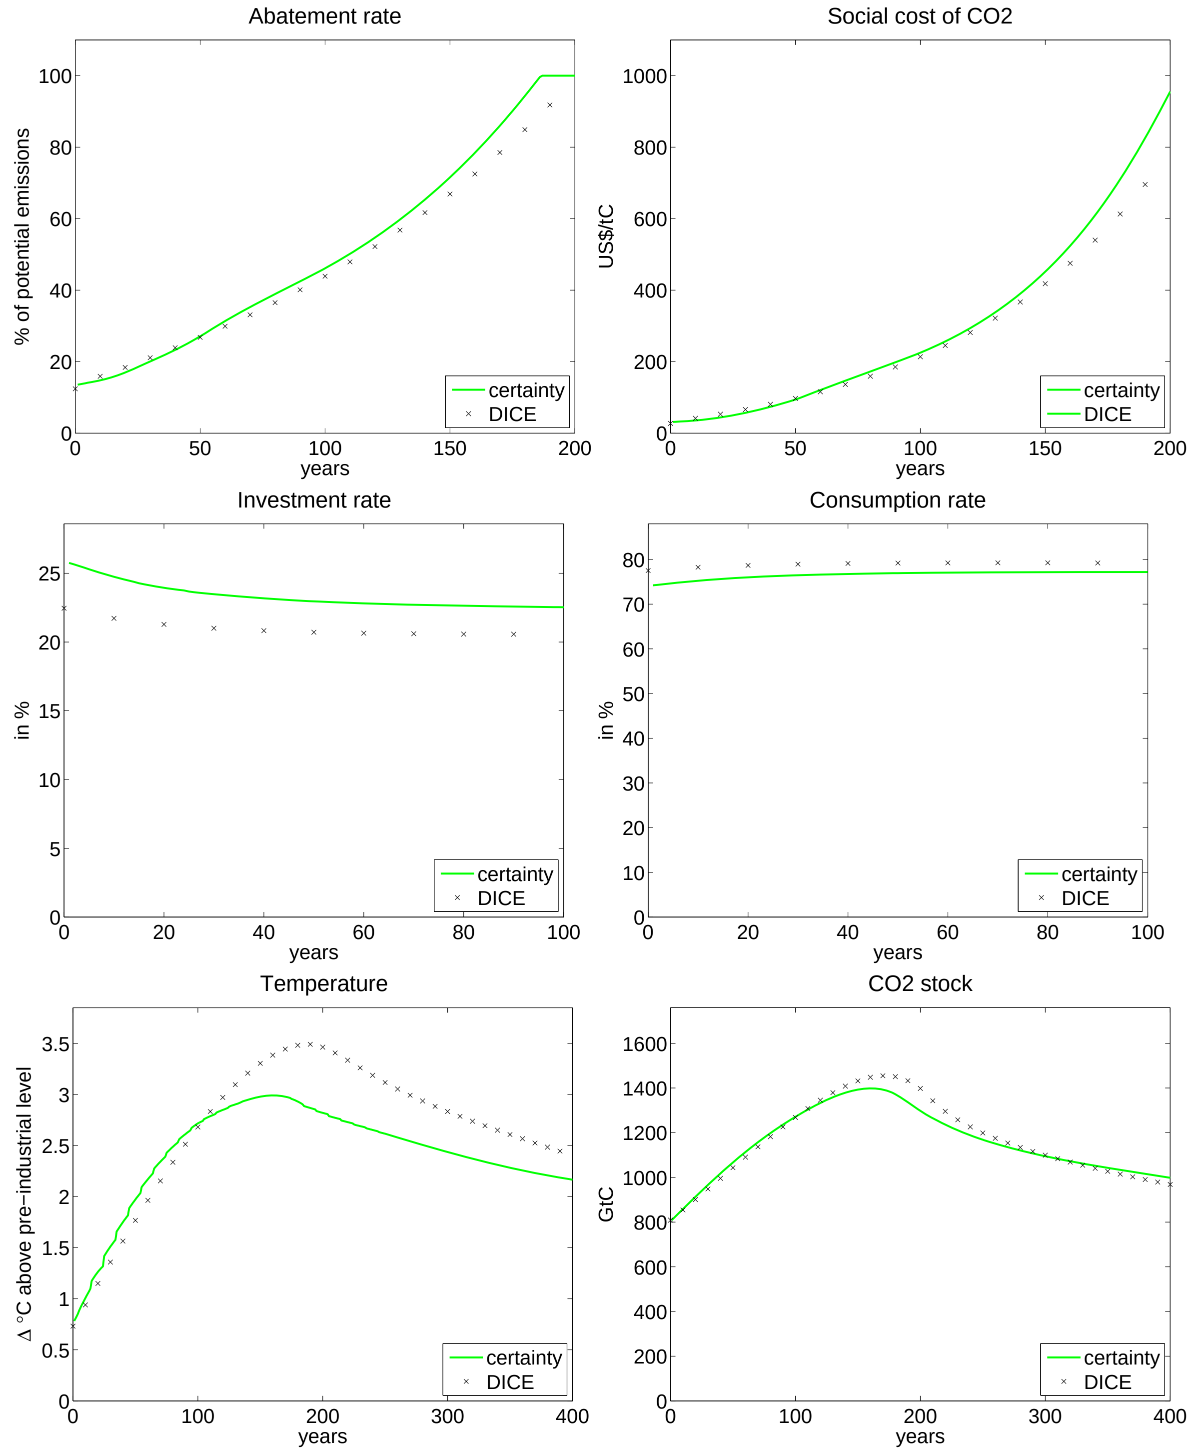

Figure 13 compares the results our recursive formulation under certainty with the original DICE model results The consumption smoothing parameter is $\eta = 2$.

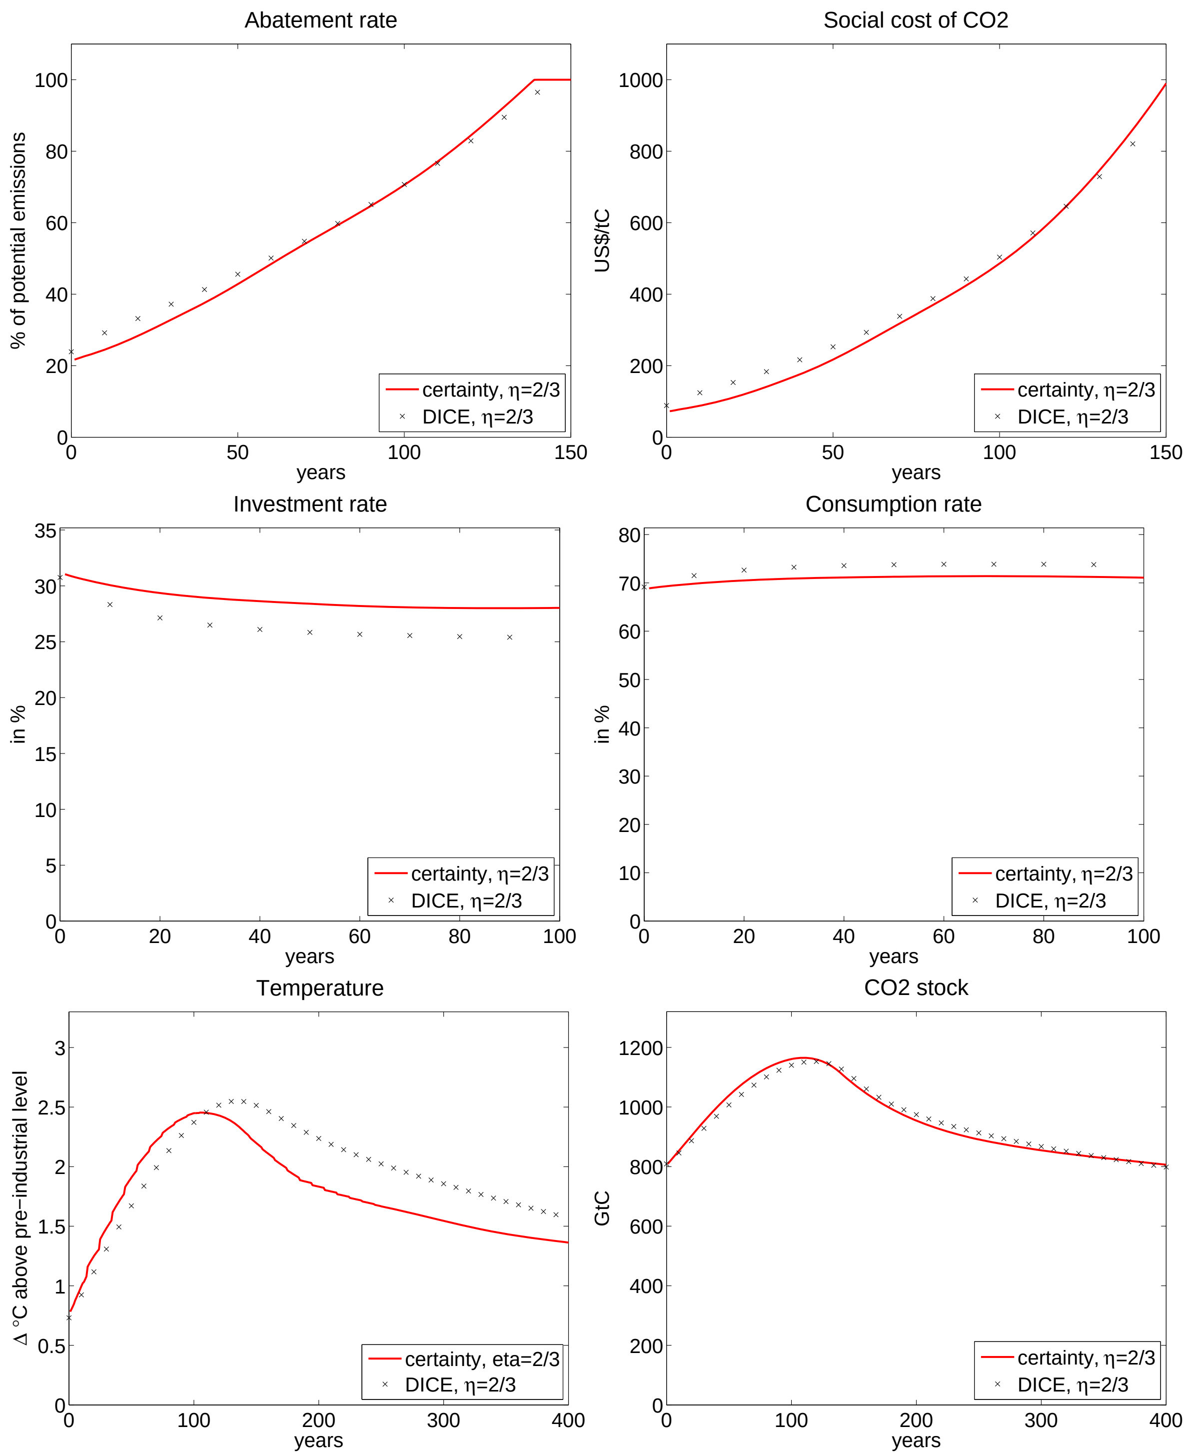

Figure 14 compares the results of our recursive formulation under certainty with the DICE model results for the same low consumption smoothing parameter $\eta = 2/3$.

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
b_2	2	damage exponent
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 in 2005; 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 0.92% used in DICE
δ_A	0.1%	rate of decline of productivity growth rate
σ_0	0.1342	CO_2 emissions 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 backup
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_{pre}	596.4	in GtC, preindustiral stock of CO2 in the atmosphere
$\delta_{M,0}$	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 2005
EF_{100}	.3	external forcing in year 2100 and beyond
f_{eql}	0.61	time invariant temperature feedback
λ_1	5.35	in $W\ m^{-2}$, additional radiative forcing from changing CO2 concentrations
λ_2	0.315	in °C $(W\ m^{-2})^{-1}$, temperature change per unit of radiative forcing