Optimal Climate Change Mitigation under Long-Term Growth Uncertainty: Stochastic Integrated Assessment and Analytic Findings[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 January 2014

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, an anonymous referee and the editors for helpful comments.

1 Introduction

The paper analyzes, numerically and analytically, the consequences of growth uncertainty for climate policy in a stochastic integrated assessment model. Economic growth is one of the most important determinants of optimal greenhouse gas mitigation. It determines both the future damages from today's greenhouse gas emissions and the marginal valuation of the resulting consumption loss.

Deterministic integrated assessment models 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. Today's optimal investment into mitigation depends on the interactions of the climate and the economy over the coming centuries. However, economic growth predictions for the far future are highly uncertain. 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.

Optimal policy depends on the evaluation of uncertainty. The standard discounted expected utility model equates intertemporal consumption smoothing with risk aversion, thus assuming a form of intertemporal risk neutrality (Traeger 2009, 2014). Epstein-Zin-Weil preferences disentangle intertemporal consumption smoothing from risk aversion, thereby improving 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 discounting and risk attitude are essential for the assessment of climate policy under uncertainty. The separation also allows us to individually identify the effects of consumption smoothing and risk aversion. In this comprehensive asset pricing framework, growth uncertainty has a strong impact on the optimal carbon tax. In contrast, a discounted expected utility evaluation implies only minimal policy adjustments. Our quantitative assessment employs Nordhaus's (2008) DICE model, used in the US federal social cost of carbon assessment Interagency Working Group on Social Cost of Carbon (2010, 2013). The disentangled risk attitude estimates in the finance literature, in combination with DICE's consumption smoothing preference, *increase* the optimal carbon tax between 20% and 45%, depending on whether shocks are iid or partially persistent. Incorporating as well the corresponding finance literature's lower estimate of consumption smoothing switches the sign of the uncertainty effect: growth uncertainty *decreases* the optimal carbon tax between 15% and 30%.

We explain this uncertainty response in the Epstein-Zin-Weil framework through an endogenous pessimism weight. In particular, we explain why investment into the climate asset decreases under uncertainty if consumption smoothing η is below unity, whereas investment into produced capital always increases under uncertainty. We find similar differences between the optimal investment in produced and climate capital in the discounted expected utility model. In DICE, the sign of the growth uncertainty effect on optimal abatement flips in both cases in the neighborhood of $\eta \approx 1$. However, we show that the interpretation of η is consumption smoothing in the more comprehensive framework and prudence in the discounted expected utility standard model. The interpretation of unity is the damage elasticity to production in the former and a combination of various parameters in the latter model.

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 order approximation to stochastic analysis does not model a decision maker's optimal response to uncertainty.[2] Nordhaus (2008) addresses growth uncertainty along a business as usual trajectory employing the Monte-Carlo approach. His analysis suggests that growth uncertainty reduces the social cost of carbon. In contrast, our stochastic analysis finds that growth uncertainty slightly increases the optimal carbon tax.

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, whereas our decision model implements a moderately persistent shock in a rational, anticipated learning framework. Gollier (2002) and Traeger (2014) extend the discount rate reasoning to Epstein-Zin-Weil preferences. Gollier et al. (2000) analyze conditions on prudence under which expected information over damage uncertainty increases the prevention effort. In a similar theory model with a numeric simulation, Ha-Duong & Treich (2004) show that in a disentangled model of consumption smoothing and risk aversion an increase in risk aversion increases pollution control, whereas an increase in consumption smoothing reduces emission control. Baker & Shittu (2008) review the mostly theoretic literature on endogenous, uncertain technological improvements of climate-friendly technologies.

The original stochastic dynamic programming implementation of the DICE model goes back to Kelly & Kolstad (1999, 2001) and we use an implementation closely related to Traeger (2012). A set of stochastic integrated assessments investigates the impact of short-term fluctuations on the choice of the optimal policy instrument (Hoel & Karp 2001, 2002, Kelly 2005, Karp & Zhang 2006, Heutel 2011, Fischer & Springborn 2011). A different strand of the literature analyzes the impact of climate sensitivity uncertainty on optimal mitigation (Kelly & Kolstad 1999, Leach 2007, Jensen & Traeger 2013, Kelly & Tan 2013) and the policy impact of tipping points (Keller et al. 2004, Lemoine & Traeger 2014, Lontzek et al. 2012). Crost & Traeger (2010) consider damage uncertainty, also employing the comprehensive Epstein-Zin-Weil preference framework.

[2] Sometimes, the decision maker is assumed to respond optimally to the ex-ante draw in every individual run, however, such a simulation can result in misleading policy suggestions, see (Crost & Traeger 2013).

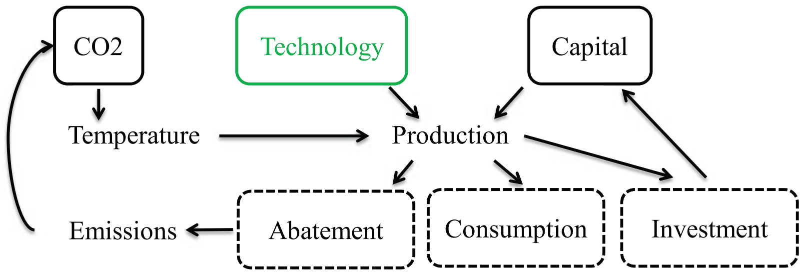

Figure 1 is an abstract representation of the climate-enriched economy model. The control variables consumption and abatement as well as the 'residual' investment are represented by dashed rectangles. The main state variables are depicted by solid rectangles. The green color indicates that the technology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

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. Global average temperature above the level of 1900 (T_t) reduces 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 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

$$Y_t = \frac{1-\Lambda_t(\mu_t)}{1+b_1T_t^2}(A_tL_t)^{1-\kappa}K_t^\kappa \ . \tag{1}$$

The damages b_1T^2 are quadratic in temperature 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. The Online Appendix III provides 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[3]

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

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

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

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

We set the standard deviation to $\sigma_z = 2.6\%$, which corresponds to twice our initial technology growth rate. This value is conservative in relation to the macroeconomic and finance literature calibrations of iid technology growth shocks. For example, Boldrin et al. (2001) use an annual standard deviation of 3.6%. Tallarini (2000) calibrates a low value of 2.3%, whereas Kaltenbrunner & Lochstoer (2010) are at the upper end of the range of values with 8.2%. The most important moment matched in those calibrations is the consumption volatility. Kocherlakota's (1996) observes for the last century of US data that the standard deviation of consumption growth is about twice its expected value, about 3.6%, which is somewhat higher than the 2.9% in the data used by Bansal & Yaron (2004).[5] We fix the mean of the growth shock so that expectations over future technology level coincide with the evolution in the certain scenario (see Online Appendix IV for detailed calculations).[6] Figure 2 illustrates the future technology level under expected growth in solid green, and the 95% (simulated) confidence interval under iid growth shocks in dashed blue. In expectation, and in the deterministic model, the productivity level of the economy increases roughly threefold over the 100 year time horizon.

In a modification, we analyze the consequences of persistence in the growth shock. While our shocks always have a persistent effect on the technology level, persistence in the growth shock implies that the growth rate itself is intertemporally correlated. Such persistent shocks are used in the macroeconomic literature to model uncertainty in growth trends (Aguiar & Gopinath 2007, Croce 2013). Here, we think of the persistent shock as a fundamental uncertain change affecting technological progress, e.g., times of economic crisis, international conflict, fundamental innovations or the

[3]Our numeric 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.

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

[5]Alternatively one can measure productivity indirectly from production and input data. For example, the Federal Reserve Bank's economic database FRED provides a total factor productivity time series, which, transformed to labor augmenting productivity, has an annualized standard deviation of 2.1%, see http://research.stlouisfed.org/fred2/series/RTFPNAUSA632NRUG.

[6]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+i} expectation equals its deterministic part for all $i > 0$ (see Online Appendix IV).

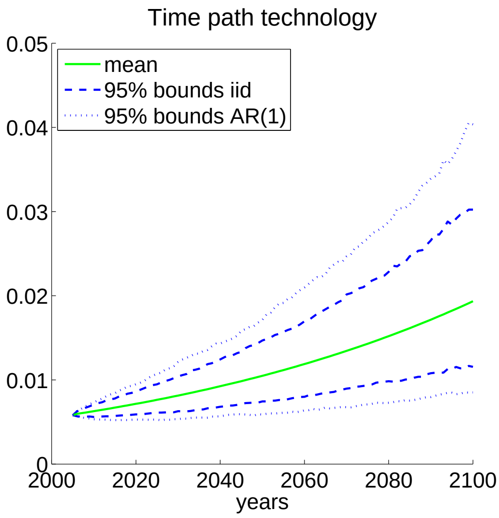

Figure 2 shows the expected draw and the 95% confidence intervals for technology time paths based on 1000 random draws of technology shock $\tilde{z}$ time paths. Both lines have an annual one-year ahead volatility of $\sigma_{\tilde{z}} = 2.6\%$. Whereas the growth rate shocks are iid for the dashed lines, the dotted lines combine an iid and a persistent AR(1) growth shock.

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

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

Choosing the standard deviations $\sigma_x = \sigma_\epsilon = 1.9\%$ results in a the same standard deviation of the overall shock $\tilde{z}_t$. 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 $w_t = 0$.[7] Our simulations assume that 50% of the ϵ-shock carries over to the growth rate of the next year: $\zeta = 0.5$. The dotted lines in Figure 2 represent the 95% (simulated) confidence interval for the technology levels over the next 100 years under such persistent growth shocks. The long term volatility is considerably higher than with an iid shock.

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

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

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 w_t captures the persistent part of last period's shock that carries over to the current period. We first state the Bellman equation for standard preferences, i.e., the time additive expected utility model:

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

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 the Online Appendix III. 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 risk-free 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

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

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 $\eta = \text{RRA}$ 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 $\text{RRA} = 10$. 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 numeric implementation for several reasons. First, modeling a random walk without mean reversion is a numeric challenge and the normalized Bellman equation introduced below converges significantly better. Second, the support of the non-normalized capital and the absolute technology level grow without bounds, limiting the planning horizon as well as the node density of a numeric 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 renormalizations, we restate the Bellman equation (5) as

$$V^*(k_\tau, M_\tau, a_\tau, \tau, w_\tau) = \max_{c_\tau, \mu_\tau} u(c_\tau) + \beta_\tau \times \tag{6}$$
$$f^{-1} \left(\mathbf{E} \left[f \left(V^*(k_{\tau+\Delta\tau}, M_{\tau+\Delta\tau}, \tilde{a}_{\tau+\Delta\tau}, \tau + \Delta\tau, \tilde{w}_{\tau+\Delta\tau}) \right) \right] \right),$$

where $u(c) = \frac{c^{1-\eta}}{1-\eta}$ and $f(v) = ((1-\eta)v)^{\frac{1-\text{RRA}}{1-\eta}}, v \in \mathbb{R}, (1-\eta) > 0$. We introduce general functions 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. The Online Appendix II spells out the detailed derivation of equation (5) 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.

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 solid 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

[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

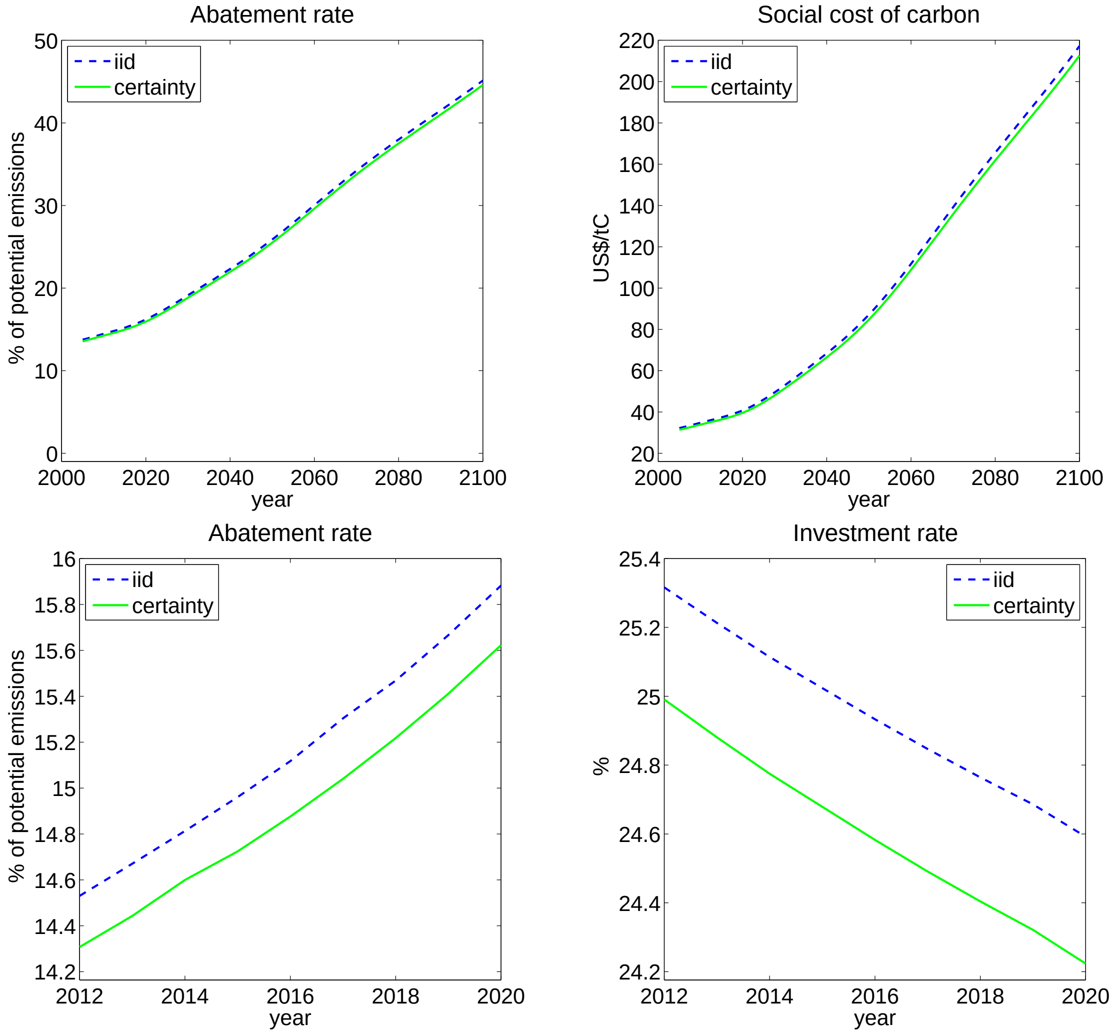

Figure 3 compares the optimal abatement rate, the social cost of carbon, and the investment rate under (iid) uncertainty to their deterministic values (standard preferences, $\eta = 2, \text{RRA} = 2$). The upper panels show the abatement rate and the social cost of carbon for the current century, the 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.

growth implies a very minor increase in optimal mitigation and the corresponding carbon tax. For the current century, the optimal abatement is 0.2-0.6 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 0.35 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 (2014) explains the smallness of the precautionary effect by pointing out that decision-makers with entangled preferences are intertemporal risk neutral (section 2.3).

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 12 in the Online Appendix I.

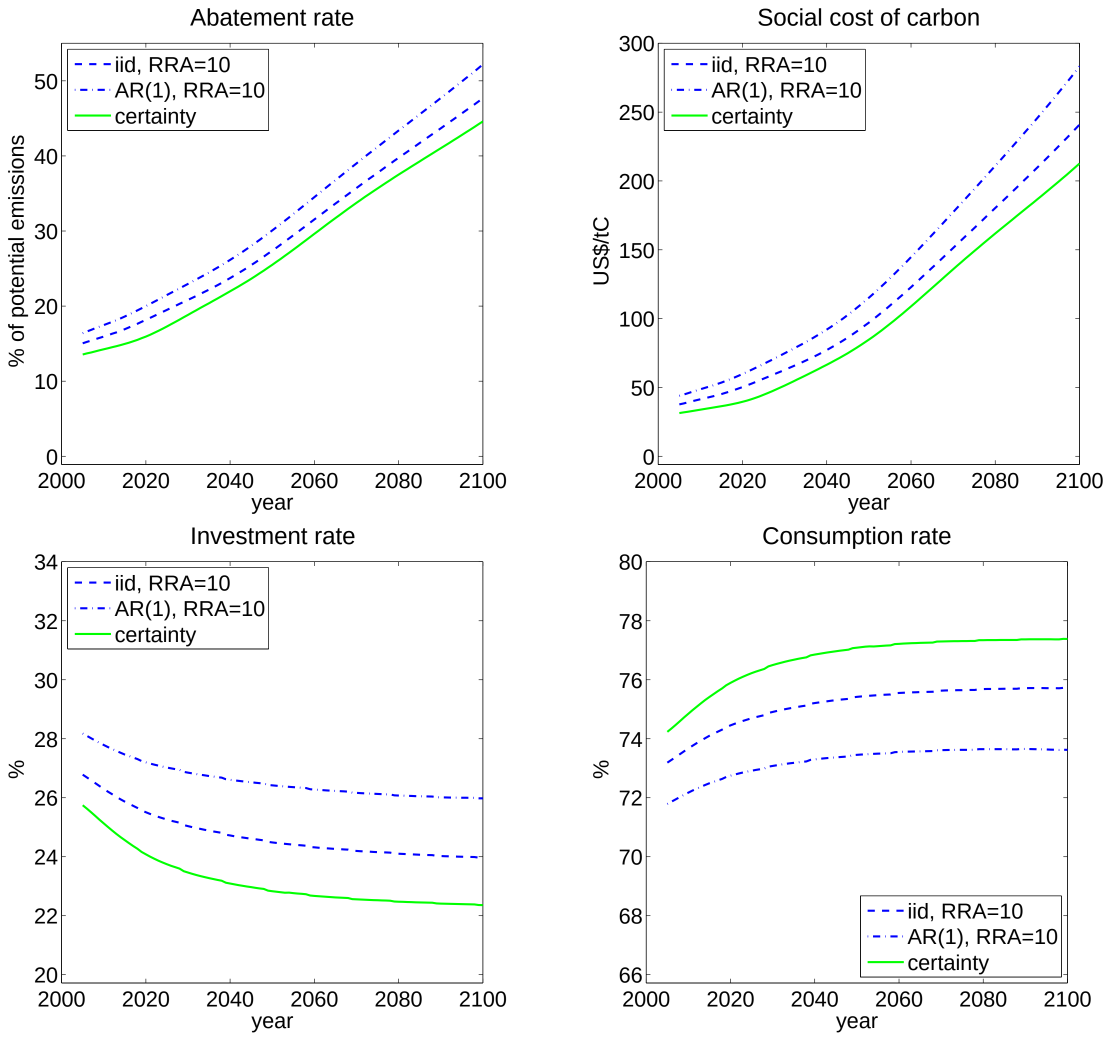

Figure 4 compares abatement, the social cost of carbon, investment rate, and consumption rate for three scenarios: certainty, an iid shock, and a persistent shock. The consumption smoothing coefficient is $\eta = 2$ and the coefficient of relative risk aversion is RRA $= 10$. Uncertainty increases abatement,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

3.2 Increasing Risk Aversion (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 RRA $= 10$. For now we keep the consumption smoothing parameter at the DICE value ($\eta = 2$).

Figure 4 shows the resulting optimal climate policy. We observe an increase in optimal 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. Note that 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 in comparison 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 $\eta = 2/3$ (see 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 (green) to 2/3 (red) 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 in Figure 5 represent optimal policy under growth uncertainty, when $\eta = 2/3$, $RRA = 10$. The optimal abatement and the social cost of carbon fall in all periods. The sign of the uncertainty effect is opposite to the one observed in the earlier settings. The optimal carbon tax decreases by 15% to $71 for the iid shock and by 32% to $56 in the case of persistentence. The abatement rate in 2012 decreases

[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.

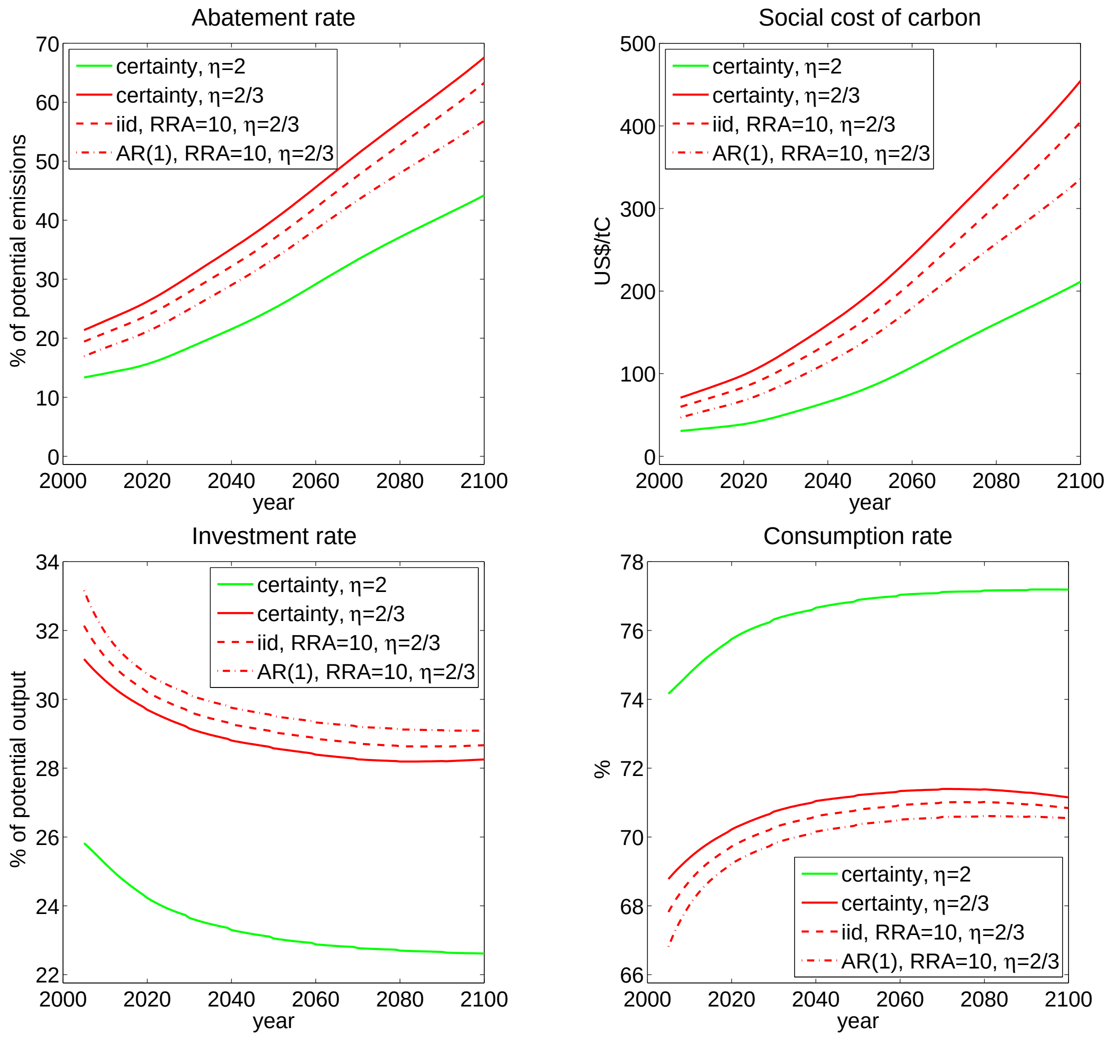

Figure 5 summarizes the results of full preference disentanglement. The solid lines depict certainty scenarios for $\eta = 2/3$ (red/dark) and $\eta = 2$ (green/light). The dashed line represents an iid shock scenario with $\eta = 2/3$, RRA $= 10$, the dashed-dotted line a persistent shock with the same 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.

by 9% to 21 percentage points for the iid shock and by 20% to 19 percentage points for the persistent shock. In contrast, investment in man-made capital still increases. The investment rate goes up by 2% for the iid shock (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 $\eta = 2$. The policy difference between the two scenarios is, however, significantly smaller under uncertainty as compared to the deterministic case.

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

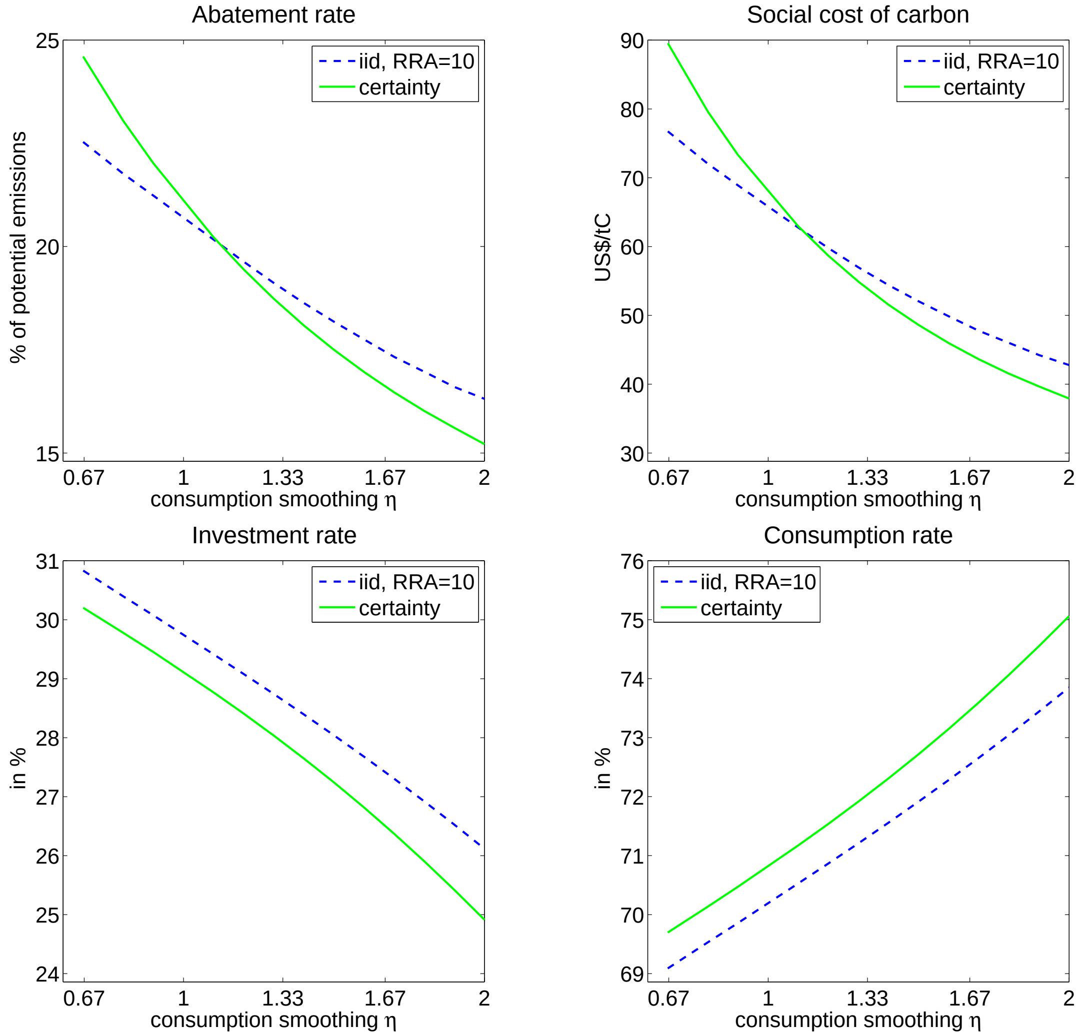

Figure 6 compares the optimal present day controls under certainty and (iid) uncertainty as a function of the consumption smoothing parameter η. Relative risk aversion is RRA $= 10$. For a high value of η, uncertainty decreases the social cost of carbon and (and vice versa). The effect switches sign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

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.

4 Analytic Discussion

This section develops an analytic understanding of the policy impact of growth uncertainty observed in section 3. We identify the structural assumptions of the DICE model driving our results. Our analysis builds on the first order conditions for optimal abatement in a generic integrated assessment model that shares DICE's model structure but uses general production, damage, and utility functions. In the first subsection, we state this first order condition for optimal abatement, which introduces the social cost of carbon along the optimal trajectory of climate and the economy. In the second subsection, we discuss the interaction between growth, climate damages, and economic valuation in a deterministic context. The deterministic result paves the way for understanding the more complex uncertainty response, in particular, in the Epstein-Zin-Weil framework. The third subsection explains how uncertain *production* shocks affect optimal abatement in the discounted expected utility model. Appendix A.3 extends the analysis to general forms of technology shocks. The forth and final subsection explains the policy response to growth uncertainty in the comprehensive Epstein-Zin-Weil framework.

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 $i > t$ as a consequence of a unit reduction of emissions in period t as $-\frac{\partial M_i}{\partial E_t}$.[10] Changing the period i carbon level affects output in that period as $\frac{\partial y_i}{\partial M_i}$. This consumption loss[11] is valued according to its marginal period i welfare change $u'(c_i)$. We translate this welfare change into period t consumption units dividing it by $u'(c_t)$, the marginal consumption value in period t. Under certainty, the social benefit in period $i > t$ from decreasing carbon in period t by one unit is therefore the product $-\frac{u'(c_i)}{u'(c_t)}\frac{\partial y_i}{\partial M_i}\frac{\partial M_i}{\partial E_t}$. The total benefit of the emission reduction in period t is the discounted sum of these benefits in all future periods. The optimal marginal abatement cost $\Lambda'(\mu_t)$ in period t is proportional to these benefits (see Appendix B):

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

In DICE, a one percent increase in the marginal abatement cost Λ' increases the abatement rate μ by approximately half a percent.

The expectation operator $\mathbf{E}_t^*$ takes expectations over all possible future sequences $A_{t+1}, A_{t+2}, \ldots$ (as opposed to just A_{t+1}), conditional on A_t (and the persistent shock w_t in the case of persistence). The first term under the sum, $\prod_{j=t}^{i} \beta_j \Pi_j P_j$, is a prudence- and pessimism-adjusted discount factor for Epstein-Zin-Weil preferences. The discount factor β_t discounts utility from period $t+1$ to period t units.[12] Π_j is a

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

[11]We are analyzing a marginal change along the optimal path, therefore, the marginal value of the consumption-capital unit is independent of whether the consumption unit is consumed, invested, or used for abatement.

[12]The discount factor $\beta_t = \exp[-\delta_u + g_{A,i}(1-\eta) + g_{L,i}]$ includes a time index because it adjusts for the time dependent labor and expected technology growth. It results from our normalization to effective labor units, which also eliminates the population weight from the Bellman equation. See Online Appendix II for details.

prudence and P_j is a pessimism adjustment. These adjustments arise endogenously in the Epstein-Zin-Weil preference framework and we discuss them in section 4.4. Under certainty and in the discounted expected utility model we have $\Pi_j = P_j = 1$.

The subsequent analysis employs equation (7) to sign the effects of growth and growth uncertainty on optimal abatement policy and to understand the driving structural assumptions of integrated assessment models. For analytic tractability, we assume a constant consumption rate.[13] Appendix A.1 discusses the minor changes resulting from endogenizing the consumption rate. The cost of carbon contribution of an individual period (summand) on the right of equation (7) responds to growth shocks as follows. First, the summand in period i responds directly to a positive growth shock in period i by affecting production and consumption in period i. Consumption in the valuation term $\frac{u'(c_i)}{u'(c_t)}$ responds proportionally to these production shocks, decreasing marginal valuation of the damage. Production also affects the level of the marginal damage $-\frac{\partial y_i}{\partial M_i}$ in period i because damages in DICE are proportional to output. More generally, we define the marginal damage function $d(y) = -\frac{\partial y_i}{\partial M_i}$ that captures the output dependence of marginal damages.

Second, the cost contribution (summand) in period i responds to growth shocks in earlier periods $j < i$. Growth rate shocks have a persistent impact on technology and production level and the shocks in $j < i$ still affect the cost of carbon contribution in period i through production and consumption increases similar to the shocks in period i. For more persistent growth shocks, a higher intertemporal correlation between shocks in periods i and j amplifies the magnitude of the uncertainty effects, which explains our findings in Figures 4 and 5. In addition to the direct output channel, a growth shock in $j < i$ also affects emissions in the earlier periods that have a delayed impact on the damage level in period i through temperature increase. The functional form of this second channel is significantly harder to determine because it has to account for the full climate dynamics. Our results show that the direct output and consumption effects, based on shocks in period i affecting the cost of carbon contribution in period i, already explain the qualitative policy impact of the growth shocks. Appendix A.2 shows that, indeed, the contribution of the emission growth effect is small as compared to the terms signing the uncertainty effect in our upcoming discussion.[14]

[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 generally not be constant, but the assumption allows us to flesh out the basic mechanisms determining abatement under certainty, under risk, and under Epstein-Zin-Weil preferences.

[14]The abstract relations derived below remain correct for the more complex delayed shock impact when interpreting the damage elasticities as capturing not only the direct output dependence, but also the emission-driven output dependence. However, the function form of this second channel is more complicated as it depends on how emission-induced radiative forcing translates into temperature change. We derive this relation explicitly in Appendix A.2.

4.2 Growth, Climate Damages, and Economic Valuation under Certainty

Ceteris paribus, a positive growth shock increases economic production in all subsequent periods, affecting future damages as well as their marginal valuation by a richer population. Using equation (7) we analyze the climate policy impact of such a growth shock, focusing on the effect of the shock within a given period.

We define the elasticity of marginal damages with respect to production as

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

In the DICE model, the marginal damage $d(y) = -\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}$ is linear in production y, and a positive production shock proportionally increases damages (equation 1).[15] Thus, the DICE model assumes $\text{Dam}_1 = 1$. We define aversion to intertemporal substitution as

$$\text{AIS}(u, c) = -\frac{u''(c)}{u'(c)} c \ .$$

In DICE, the isoelstic utility function implies $\text{AIS} = \eta$, and Nordhaus (2008) assumes $\eta = 2$.

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

$$\text{AIS}(u, c) < \text{Dam}_1(d, y) \qquad \Leftrightarrow \qquad \eta < 1 \quad \text{in DICE} \ . \tag{8}$$

Technological growth increases the social cost of carbon and the optimal abatement rate if (and only if) the wealth-induced decline in marginal valuation of a unit damage, characterized by aversion to intertemporal substitution $\text{AIS}(u, c)$, is dominated by the increase in physical damages captured by $\text{Dam}_1(d, y)$. The finding that, under positive growth, a lower aversion to intertemporal substitution increases the cost of carbon is often stated in terms of the social discount rate: lowering the aversion to intertemporal substitution lowers the consumption discount rate and, thus, increases the weight given to future damages. In addition, a higher sensitivity of (marginal) damages to production strengthens optimal abatement.

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 13 in the Online Appendix I. From equation (8) we also learn that logarithmic utility in a DICE-like model with linear-in-production damages renders climate policy independent of the technology (and production) level. This case, where $\eta = 1$ and $\text{Dam}_1(d, y) = 1$, is the setting of Golosov et al. (2011) analytic integrated assessment model.

[15] $\frac{\partial y_{\tau+1}}{\partial M_{\tau+1}}|_{k,M,T} = -g(M, T, t)\, y_t$ where $g(M, T, t)$ depends on the states of the climate system only. See Appendix A.2 for the small, indirect impact of earlier production shocks on the damage level through the emission growth channel.

4.3 The Uncertainty Effect in the Discounted Expected Utility Standard Model

In the real world, growth shocks are both positive and negative, and the asymmetry in the damage response to positive and negative shocks determines optimal mitigation policy. In this section, we analyze the consequences of a mean-zero shock on production. Appendix A.3 extends the result to general forms of technological progress, and we briefly summarize the findings at the end of this section. We use the discounted expected utility standard model for uncertainty evaluation and, therefore, keep $\Pi_j = Pi_j = 1$ in equation (7).

Jensen's inequality implies that a mean-zero production shock in period i raises the period i contribution to the cost of carbon, if the product $u'(c)d(y)$ on the right hand side of (7) is convex. Therefore, we introduce elasticities characterizing the second order changes of marginal utility and damages

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

which characterizes the convexity of marginal damage in the production level. The DICE model assumes that damages are linear in production so that $\text{Dam}_2(d, y) = 0$. Similarly, we define Kimball's (1990) measure of relative prudence

$$\text{Prud}(u, c) = -\frac{u'''(c)}{u''(c)} c \ .$$

The measure is known to characterize the precautionary savings (investment) response to income uncertainty. For the isoelastic utility function in DICE we have Prud $= 1 + \eta = 3$. The positivity of relative prudence explains the increase of investment in produced capital under uncertainty. We note that in the discounted expected utility model $\text{AIS}(u, c) = \text{RRA}$ ($= \eta = 2$ in DICE) jointly characterizes aversion to intertemporal substitution and relative risk aversion.

Appendix B shows that production uncertainty increases optimal abatement if

$$\begin{aligned} \text{Prud}(u, c) &> 2\ \text{Dam}_1(d, y) \ - \frac{\text{Dam}_1(d, y)}{\text{AIS}(u, c)}\ \text{Dam}_2(d, y) \qquad (9) \\ \Leftrightarrow \quad 1 + \eta \ &> 2 \ * \ 1 \ - \ 0 \qquad \text{in DICE} \ . \end{aligned}$$

In contrast to the uncertainty response of investment into produced capital, abatement only increases if prudence is not only positive, but also dominates the sensitivity of damages to production shocks. In the DICE model, damages are linear in production and utility is isoelastic. Then, production shocks increase optimal abatement if $\eta > 1$. We emphasize that the interpretation of this equation and the driving force for increased abatement is neither risk aversion nor consumption smoothing dominating unity, but prudence dominating the damage elasticity.

Interpreting equation (9), positive prudence on the left hand side characterizes a decision maker whose valuation of marginal consumption reacts more strongly to negative than to positive growth shocks: she increases the value of the physical damage under low growth more than she lowers the valuation of the physical damage under high growth. This asymmetry in the response would make the prudent decision maker strengthen his mitigation effort under uncertainty, if the physical damages were insensitive to growth shocks. However, the physical damages are lower in the low growth corresponding to a higher marginal valuation. Therefore, the optimal abatement response to uncertainty is positive only if the prudence effect dominates (twice) the damage sensitivity to production (first term on the right of equation 9). An additional contribution increases optimal abatement if damages are convex in the production level: damages under high growth increase more than they diminish under low growth, which increases the cost of carbon along the optimal path and, thus, optimal abatement (second term on the right of equation 9). This latter effect is stronger whenever a positive growth shock increases abatement in the first place (damage sensitivity dominates aversion to intertemporal substitution).

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 ($\text{Dam}_1 = \text{Dam}_2 = 0$). 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, the more convex damages reduce the requirements on prudence. For isoelastic preferences, however, prudence ($= 1+\eta$) and risk aversion (=consumption smoothing=η) are dependent, and lowering prudence on the left hand side also reduces the right hand side of equation (9) by reducing risk aversion. For the example of quadratic-in-production damages, the η-region where growth uncertainty decreases optimal abatement shifts from $\eta \in [0, 1]$ to the interval $\eta \in [1, 2]$.[16]

Appendix A.3 extends the above analysis to general forms of technological progress. Uncertain technological progress implies mean zero shocks on production only in the case of total factor productivity enhancing technological progress. Our labor augmenting technological progress, for example, introduces an additional non-linearity between shocks and damage valuation. We show that, in the special case of DICE, the condition for uncertainty to increase abatement stays $\eta > 1$ for labor augmenting technological progress. More generally, however, the convexity in the relation between technological progress and production interacts with both the elasticity of marginal damages with respect to production and the aversion to intertemporal substitution. In the case of quadratic-in-production damages, for example, the η-range where uncertainty reduces optimal mitigation enlarges from $\eta \in [1, 2]$ under mean-zero total factor productivity shocks to $\eta \in [0.6, 2]$ for mean-zero shocks on labor augmenting technology (assuming DICE's $\kappa = 0.3$).

[16] For quadratic damages in production we find $\text{Dam}_1 = 2$ and $\text{Dam}_2 = 1$ so that equation (9) is positive if and only if $(\eta - 1)(\eta - 2) > 1$ and, thus, the uncertainty effect on abatement is negative if $1 < \eta < 2$.

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, and a sign switch in the parameter η. We characterize the corresponding adjustments to the social cost of carbon in equation (7) using a prudence factor Π_j and a pessimism factor P_j. We now explain these factors and discuss how they modify optimal climate policy under uncertainty.

4.4.1 Precautionary savings in the Epstein-Zin-Weil model

In Appendix B we show that the first order condition for consumption optimization implies[17]

$$u'(c_t) \propto \Pi_t \, \mathbf{E}_t \, P_t \, \frac{\partial V_{t+1}}{\partial k_{t+1}} \,. \tag{10}$$

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 (10) increases optimal marginal utility of the last consumption unit and, thus, decreases the consumption level and increases investment (savings).

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 (10) and, thus, investment under uncertainty if, and only if, absolute intertemporal risk aversion $-\frac{f''}{f'}$ decreases in welfare (Traeger 2011). We can rewrite the condition of decreasing absolute intertemporal risk aversion as $\text{Prud}(f, V) > \text{RRA}(f, V)$, i.e., prudence (of f evaluated at V) dominating (intertemporal) risk aversion. This condition is always met for Epstein-Zin-Weil preferences due to their isoelastic form, and it motivates naming Π a prudence term. However, a mean-zero technology shock does not necessarily produce mean-zero welfare shocks: an additional consumption is unit is valued higher in the case of a negative shock than in the case of a richer future. 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 (see Figure 10 in Appendix A.4). For the $\eta = 2$ scenario, however, we find that the value function is strongly concave in a, biasing down the expected value of V. In consequence, investment into future, produced capital is less attractive in

[17] The proportionality absorbs exogenous terms that do not change under uncertainty or with the preference specification.

the scenario with faster falling marginal utility. In the $\eta = 2$ scenario, this strong decreases in marginal utility of the richer future generations under a positive growth shock dominates the prudence effect in Π_t and implies that overall the term slightly decreases optimal investment. The resulting uncertainty corrections are relatively small and dominated by the pessimism effect discussed in the next paragraph.

The quantitatively dominating uncertainty impact on consumption, investment, and abatement 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 10). The pessimism effect, therefore, increases investment (savings) and decreases consumption.

4.4.2 Abatement in the Epstein-Zin-Weil model

In Appendix B we derive the following first order condition for marginal expenditure on abatement as a fraction of total production:[18]

$$\Lambda'(\mu_t) \propto \frac{\mathbf{E}_t \ P_t \ \left(-\frac{\partial V_{t+1}}{\partial M_{t+1}}\right)}{\mathbf{E}_t \ P_t \ \frac{\partial V_{t+1}}{\partial k_{t+1}}} . \tag{11}$$

Under certainty, equation (11) states that the optimal abatement rate increases in the marginal value of climate capital (deteriorating in M, hence $-\frac{\partial V_{t+1}}{\partial M_{t+1}}$) and decreases in the marginal value of produced capital (opportunity cost).

The prudence term Π_t cancels in equation (11), as it equally affects the marginal value of produced capital and climate capital. The denominator on the right hand side of equation (11) measures the pessimism weighted marginal value of capital. We discussed in section 4.4.1 that this pessimism weighted capital value increases under uncertainty. In equation (11) it therefore reduces the optimal abatement rate by increasing the opportunity value of investing in produced capital.

In contrast to the marginal value of produced capital, the response of the marginal value of climate capital $-\frac{\partial V_{t+1}}{\partial M_{t+1}} (> 0)$ to growth shocks is ambiguous. As we discussed in section 4.2, the value of an emission reduction increases under a positive growth shock if and only if damages are more sensitive to production shocks than the marginal value of consumption: $\text{AIS}(u, c) < \text{Dam}_1(d, y)$. Figure 11 in Appendix A.4 shows the

[18]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.

immediate implication for the marginal value of climate capital $-\frac{\partial V_{t+1}}{\partial M_{t+1}}$ as a function of the (normalized) technology level a: for $\eta = 2$ the marginal value of climate capital decreases in the technology level, whereas 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 the presence of growth uncertainty.

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, uncertainty significantly increases optimal abatement. In the $\eta = \frac{2}{3}$ scenario, the bad states of the world corresponding to low technology realizations 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 (opportunity) value of produced capital. Thus, the pessimism weighting reduces the optimal abatement rate in the $\eta = \frac{2}{3}$ scenario.

We close this section relating the abatement effect directly to our formula for the social cost of carbon in equation (7). Again, Jensen's inequality is the key to understanding the uncertainty effect on optimal abatement. In addition to the contributions $u'(c(a))$ and $d(y(a))$ whose joint convexity we discussed in section 4.3, we now have to account for the prudence term $\Pi_t(V(a))$, and the pessimism term $P_t(V(a))$. We limit our discussion to the dominant contribution to the product's convexity: $\frac{d}{da}P_t(V(a))\frac{d}{da}\Big(u'(c(a))d(y(a))\Big)$. 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. The derivative $\frac{d}{da}(u'(\cdot)d(\cdot))$ signs the abatement response to a deterministic growth increase (see section 4.2 and Appendix B). Thus, the Epstein-Zin-Weil framework gives rise to a dominant uncertainty contribution that has the same effect on optimal abatement policy as a growth reduction. The main message, corresponding to the quantitatively dominating uncertainty contribution, is simple: uncertainty in growth has the same policy effect as a deterministic growth decrease, independent of whether growth increases or decreases abatement policy (condition 8).

Variations of the above insight are as follows. Fixing a positive growth rate and the damage function, an increase in aversion to intertemporal substitution causes a decrease in optimal abatement through the deterministic channel. At the same time, this increase is partially offset by uncertainty in the Epstein-Zin-Weil framework, as observed in Figure 6. Similarly, fixing a positive growth rate and the utility function, an increase in the damage sensitivity to production increases optimal abatement through the deterministic channel, but uncertainty partially offsets this policy strengthening: the pessimism term puts more weight on the low growth states corresponding to the lower damages.

5 Conclusions

We quantify and explain the consequences of growth uncertainty for optimal mitigation policy in a stochastic dynamic programming model based on Nordhaus's (2008) integrated assessment model DICE. We identify the structural assumptions that drive the results. Deterministic growth increases optimal abatement if the sensitivity of marginal damages to production outweighs the aversion to intertemporal substitution, i.e., the higher physical damages in the future outweigh the wealth-driven reduction in marginal valuation. In the standard DICE model specification, deterministic growth decreases the optimal present-day carbon tax.

In the real world, growth is uncertain and the asymmetry in the damage response to positive and negative shocks determines optimal mitigation policy. A prudent decision maker's valuation of marginal consumption reacts more strongly to negative than to positive growth shocks: he increases the valuation of the physical damage under low growth more than he lowers the valuation of the physical damage under high growth. At the same time, physical damages are lower in the low growth world, where damages receive the higher marginal valuation. In total, the optimal abatement response to uncertainty is positive only if the prudence effect dominates (twice) the damage sensitivity to production. In the standard DICE specification, this condition is satisfied and, numerically, we find a small increase of the optimal abatement rate under uncertainty. This finding contrasts with an earlier Monte-Carlo based study by Nordhaus (2008) that suggested a reduction of the social cost of carbon under growth uncertainty.

Optimal climate policy responds much stronger to uncertainty under a more comprehensive valuation of uncertainty using Epstein-Zin-Weil preferences. We characterized this valuation change by means of a prudence and a pessimism weighting effect. Quantitatively, the dominating pessimism effect increases the effective weight of the bad states of the world with low growth and poorer future generations. Optimal abatement increases under uncertainty if these low growth realizations increase the marginal valuation of the climate asset (the social cost of carbon), which holds if consumption smoothing dominates the sensitivity of damages to production. This condition coincides with the condition that a deterministic growth reduction increases optimal abatement. Thus, uncertainty in the Epstein-Zin-Weil framework always has the opposite policy impact of deterministic growth. Quantitatively, increasing the coefficient of relative risk aversion to its disentangled estimate from the finance literature *increases* 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 persistent shock (to about \$50). Here, we are in a model where deterministic growth reduces optimal abatement and uncertainty brings abatement back up.

Lowering the consumption smoothing parameter to the disentangled estimates of the empirical asset pricing literature, which explains the low risk-free discount rate, flips the sign of the uncertainty effect. Then, we are in a model where deterministic growth increases optimal abatement because the production shocks increase the physical damages more than the wealth increase reduces their marginal valuation. Hence, under certainty, we have a much stronger abatement policy. Growth uncertainty then *reduces* the optimal carbon tax by 15% under iid shocks (to slightly above \$70) and by 30% under shock persistence (to below \$60). These uncertainty adjustments of the optimal carbon tax are larger than the risk premia found in stochastic integrated assessments of damage uncertainty in Crost & Traeger (2010), and of the same order of magnitude as the adjustments resulting from the possibility of stochastic carbon cycle and feedback tipping points analyzed in Lemoine & Traeger (2014).

We conclude that all of our empirical simulations give rise to a higher optimal carbon tax under growth uncertainty than does the base case deterministic DICE model. Quantitatively, this increase is significant when disentangling risk aversion and risk premia from intertemporal consumption smoothing and the risk-free discount rate. The resulting higher carbon tax is less sensitive to the consumption smoothing parameter than in the standard model because the dominating uncertainty contribution in the Epstein-Zin-Weil model counteracts the deterministic response to consumption smoothing.

References

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

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

Anthoff, D. & Tol, R. 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.

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.

Boldrin, M., Christiano, L. J. & Fisher, J. D. M. (2001), 'Habit persistence, asset returns, and the business cycle', *The American Economic Review* **91**(1), pp. 149–166.

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

Croce, M. M. (2013), Long-run productivity risk: A new hope for production-based asset pricing?, University of north carolina working paper.

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.

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.

Gollier, C. (2002), 'Discounting an uncertain future', *Journal of Public Economics* **85**, 149–166.

Gollier, C., Jullien, B. & Treich, N. (2000), 'Scientific progress and irreversibility: An economic interpretation of the 'precautionary principle'', *Journal of Public Economics* **75**, 229–253.

Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2011), 'Optimal taxes on fossil fuel in general equilibrium', *NBER Working Paper* (17348).

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

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

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

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

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* .

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

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.

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

Keller, K., Bolker, B. M. & Bradford, D. F. (2004), 'Uncertain climate thresholds and optimal economic growth', *Journal of Environmental Economics and Management* **48**, 723–741.

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

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

Kelly, D. L. & Kolstad, C. D. (2001), 'Solving infinite horizon growth models with an environmental sector', *Computational Economics* **18**, 217–231.

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

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

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

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

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

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

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

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

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* .

Tallarini, T. D. (2000), 'Risk-sensitive real business cycles', *Journal of Monetary Economics* **45**(3), 507 – 532.

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. (2014), '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.

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.

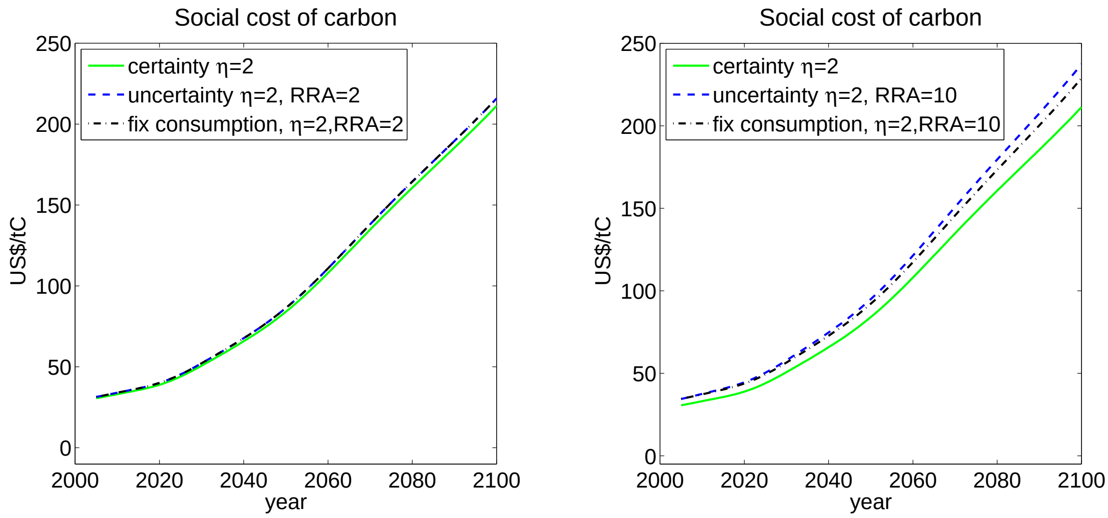

Figure 7 shows the optimal carbon tax under certainty, uncertainty, and uncertainty with the consumption rate fixed at the deterministically optimal level. The left panel displays standard entangled preferences (η = RRA = 2), the right panel shows the disentangled preference scenario ($\eta = 2$ and 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.

Appendix

A Further Results

A.1 Endogenous savings

For the purpose of analytic tractability, we assumed a constant consumption rate in section 4. However, Figure 6 in section 3 shows that the consumption rate decreases under uncertainty. In this section we show numerically that relaxing the constant consumption rate assumption is innocuous for our analytic results. Figure 7 shows the optimal carbon tax under certainty and the optimal carbon tax under uncertainty for the case of an optimal consumption rate and for the case where the consumption rate is fixed to the deterministic level. The left panel shows the case of the discounted expected untility model, the right panel disentangles relative risk aversion from consumption smoothing. Fixing the consumption rate to the deterministic level has an almost imperceptible effect for standard preferences, and a small but notable negative one for the case of Epstein-Zin-Weil preferences. Note that with a fixed consumption rate investment is lower (absence of precautionary savings), so overall emissions are still lower.

Figure 8 displays the consumption control rule as a function of the technology level for the discounted expected utility model. The agent's consumption rate decreases under a positive growth shock and increases under a negative growth shock. Absolute consumption still 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), acting similarly to a lowering

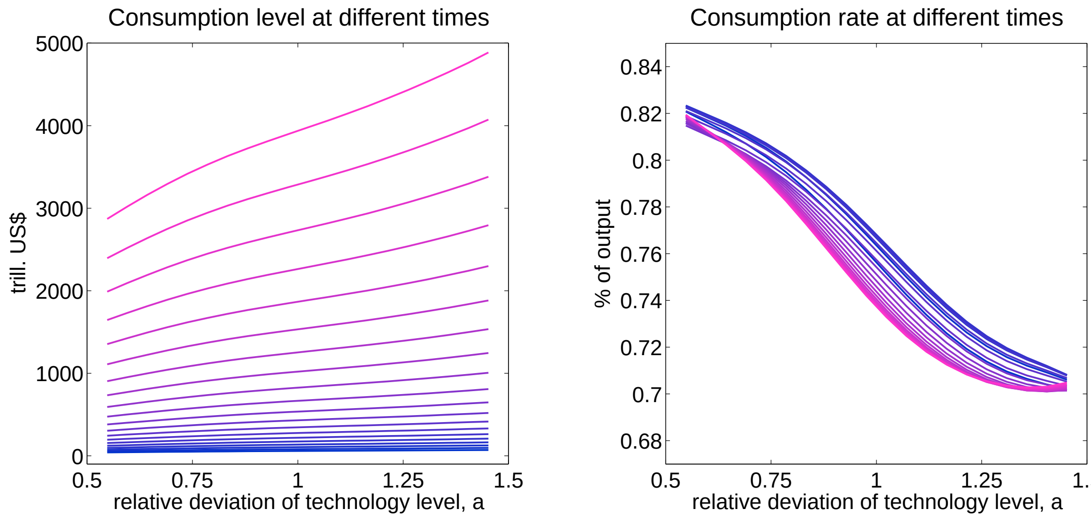

Figure 8 shows optimal consumption and the consumption rate over the normalized technology level ($a = \frac{A}{A^{det}}$) under uncertainty with entangled preferences (η = RRA = 2). The 20 different lines correspond to different points in time along the optimal path. Time intervals are evenly spaced from 2005 (dark blue) to 2400 (light pink). While the consumption level increases slightly in technology level, the consumption rate decreases.

of η. The smaller change in marginal utility implies that endogenous consumption reduces the uncertainty effect for $\eta = 2$, and increases the value of η at which the uncertainty effect flips from positive to negative.

Figure 7 also shows how the response to uncertainty of abatement and investment differ. In particular, the uncertainty adjustment of abatement does not merely mirror the changes in capital investment under uncertainty. In principle, climate change policy could be solely a response to changes in investment which then change the level of emissions. Observing that climate policy depends very little on whether capital investment is exogenous or endogenous, which is yet another way to rule out this possibility.

A.2 Analytics of the emissions growth effect

This appendix analyzes how the analytic conditions change when we analyze how shocks in a period $j < i$ affect the cost of carbon in the later period i. First, the growth shock in period j affects output in period i because (even iid) growth shocks have a persistent effect on the technology level. In addition, the growth shock in $j < i$ affects emissions in the earlier periods, which have a delayed impact on the damage level in period i through temperature increase. Therefore, we re-define the marginal damage function $\hat{d}_i(y_i, T_i) = -\frac{\partial y_i}{\partial M_i}$ taking explicit account of both growth shock channels. To simplify the resulting terms, we keep the DICE model's assumption that temperature dependence and output dependence of the damage function are multiplicatively separable: $\hat{d}_i(y_i, T_i) = g(y_i)h(T_i)$. We denote the dependence of output and temperature in period i on the technology level (and shock) in period j by $y_i(a_j)$ and $T_i(y_j(a_j))$. Then, a deterministic growth shock in period j increases the cost of carbon contribution of period i in equation (7) if

$$\begin{aligned} &\frac{\frac{d}{da_j} u'(\alpha y_i(a_j)) g(y_i(a_j)) h(T_i(a_j))}{u'(\alpha y_i(a_j)) g(y_i(a_j)) h(T_i(a_j))} > 0 \\ \Leftrightarrow\ &\frac{u''(\alpha y_i(a_j))}{u'(\alpha y_i(a_j))} \alpha \frac{\partial y_i}{\partial a_j} + \frac{g'(y_i(a_j))}{g(y_i(a_j))} \frac{\partial y_i}{\partial a_j} + \frac{h'(T_i(a_j))}{h(T_i(a_j))} \frac{\partial T_i}{\partial a_j} > 0 \\ \Leftrightarrow\ &\underbrace{\frac{u''(c_i)}{u'(c_i)} c_i}_{-\,\text{AIS}} + \underbrace{\frac{g'(y_i(a_j))}{g(y_i(a_j))}}_{\text{Dam}_1(d,y)} + \underbrace{\frac{h'(T_i(a_j))}{h(T_i(a_j))} \frac{\partial T_i}{\partial a_j} \frac{y_i}{\frac{\partial y_i}{\partial a_j}}}_{\text{Correction Term}} > 0 \ , \end{aligned}$$

where α denotes the consumption rate, which is fix by assumption. Thus, a shock in period $j < i$ increases the period i contribution to the carbon cost in equation (7) if

$$\text{AIS}(u, c) < \text{Dam}_1(d, y) + \frac{h'(T_i(a_j))}{h(T_i(a_j))} \frac{\partial T_i}{\partial a_j} \frac{y_i}{\frac{\partial y_i}{\partial a_j}} \tag{12}$$

$$\Leftrightarrow \eta < 1 + \frac{h'(T_i(a_j))}{h(T_i(a_j))} \frac{\partial T_i}{\partial a_j} \frac{y_i}{\frac{\partial y_i}{\partial a_j}} \quad \text{in DICE} \ ,$$

extending the condition in equation (8) that only considered the effect of a shock in period i on the period i contribution ($\frac{\partial T_i}{\partial a_j} = 0$ for j$\geq i$).

We can calculate the correction term for the DICE model, making use of the climate equations (and constants) for emissions, radiative forcing, temperature increase, and damages found in the supplementary online material. Note that the growth rate shock a_j increases output and, thus, emissions in all future periods. The following term only characterizes the contribution through the output and emission increase in period j itself (we therefore write $T_i(y_j(a_j))$)

$$\frac{h'(T_i(y_j(a_j)))}{h(T_i(y_j(a_j)))} \frac{\partial T_i}{\partial y_j} \frac{\partial y_j}{\partial a_j} \frac{y_i}{\frac{\partial y_i}{\partial a_j}} = \frac{b_1 b_2 T_i^{b_2 - 1}}{1 + b_1 T_i^{b_2}} \ \frac{s \chi_i}{(1-\kappa) \ln 2} \ \left[\Pi_{t=j+1}^{i-1} (1 - \delta_M(M, t))\right] \frac{E_j}{M_i}.$$

The order of magnitude of the term is mostly determined by the ratio of the emission flow over the emission stock $\frac{E_j}{M_i}$, which is of the order 10^{-2}, making the correction term small as compared to the other contributions. To get the full contribution, we have to sum up the resulting temperature level effects of the shock a_j resulting from increasing $y_j, y_{j+1}, ..., y_{i-1}$ and the corresponding emission levels $E_j, E_{j+1}, ..., E_{i-1}$. Figure 9 presents the result for a shock in the present year 2014. Given the persistence of the growth shock on the technology level, the temperature effect becomes larger the further we go into the future. It's magnitude, however, stays small as compared to the contributions of AIS and $\text{Dam}_1(d, y)$, which are of unit order. Moreover, by depicting the impact of a shock in 2014 on the future periods, we depict the largest possible contribution of the correction term relative to AIS and $\text{Dam}_1(d, y)$; later shocks will have a smaller impact relative to AIS and $\text{Dam}_1(d, y)$ in any given year. Finally, the uncertainty contribution that is signed by the relation of AIS, $\text{Dam}_1(d, y)$,

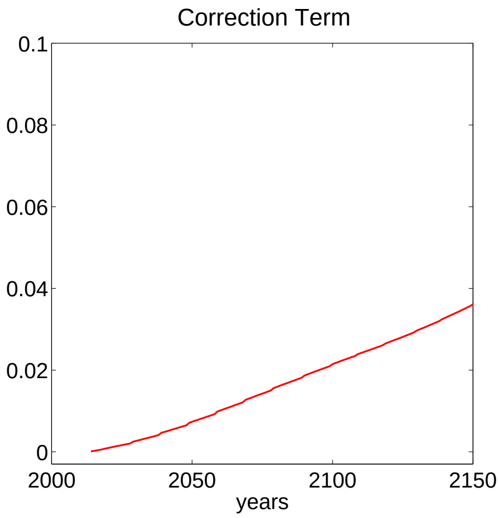

Figure 9 shows the additional term entering the condition for a positive growth shock to increase optimal abatement when the shock period precedes the evaluation year. Here, the persistent growth shock occurs in 2014. The x-axis gives the year in which we evaluate the social cost of an additional unit of CO_2 emitted in the present. The y-axis gives the normalized contribution of the correction term $\frac{h'(T_i(a_j))}{h(T_i(a_j))}\frac{\partial T_i}{\partial a_j}\frac{y_i}{\frac{\partial y_i}{\partial a_j}}$ in the inequality (12). Its size is small relative to the other contributions of the the inequality which are of unit order. Note that the graph does not show absolute contributions, the absolute contribution is discounted and reduced by emission decay for later years.

and the correction term will have less impact on the social cost of carbon the further the contribution lies in the future because of both discounting (β_t) and the decay of the additional carbon unit emitted today whose effect equation (7) measures ($\frac{\partial M_{i+1}}{\partial E_{t+1}}$.

The smallness of the correction term is mostly a consequence of the climate dynamics captured by $\frac{\partial T_i}{\partial y_j}$ that gives rise to the ratio $\frac{E_j}{M_i}$, where the term E_j results from emissions being proportional to production and $\frac{1}{M_i}$ results from radiative forcing being logarithmic in the carbon stock. Thus, the correction term will be small for most model assumptions, unless the model contains an extremely steep damage function so that the term $\frac{h'(T_i)}{h(T_i)}$ outweighs the smallness of temperature effect. Finally, note that there is one more, even smaller correction that stems from a growth impact on equation (7)'s term $\frac{\partial M_{i+1}}{\partial E_{t+1}}$. The growth-shock-triggered emissions increase the overall carbon stock, which slowly saturates some of the carbon sinks in the long run. At the same time, however, higher emissions imply a higher partial pressure in the atmosphere, which increases the emission flow into the carbon sinks. Accounting for both effects, the rate of atmospheric carbon removal is fairly insensitive to the emission level changes over the next century, and the impact of its change in emission levels is negligible in our derived conditions.

Condition 12 characterizes the general, delayed growth shock impact on the cost of carbon. The condition reappears in the Epstein-Zin-Weil context for the uncertain growth shock with opposite sign in relation to the pessimism weighting. It is one of the reasons why the abatement rate under certainty and the abatement rate under uncertainty cross where η is slightly larger than unity rather than exactly unity.

In the discounted expected utility model, the relevant condition for uncertainty to increase abatement was described by the more complex equation (9). We derived the correction term to $\mathrm{Dam}_1(d, y)$ above. In addition, a growth shock preceding the evaluation period modifies the convexity of the marginal damages characterized by the term $\mathrm{Dam}_2(d, y)$. This marginal damage convexity measure $\mathrm{Dam}_2(d, y)$ now has to account for the temperature level effect and the second order cross-derivatives of $\hat{d}_i(y_i, T_i) = g(y_i)h(T_i)$. It is straight-forward to show that the terms other than $\mathrm{Dam}_2(d, y) = \frac{g''(y)}{g'(y)}y$ contain either $\frac{\partial T_i}{\partial y_j}$, of the order $\frac{E_j}{M_i}$, or $\frac{\partial^2 T_i}{\partial y_j^2}$ of a similarly small order. Thus, also the corrections to condition (9) for a growth shock preceding the evaluation period are small as compared to the relevant contributions stated in section 4.3.

A.3 General 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 (see Appendix IV). Production is concave in labor and, thus, a shock that keeps the expected technology level constant slightly decreases expected production.[19] 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 B)

$$\mathrm{Prud}(u, c) > \mathrm{Dam}_1(d, y)\left[2 - \frac{\mathrm{Dam}_2(d, y)}{\mathrm{AIS}(u, c)}\right] + \frac{\mathrm{Tech}_2(\mathrm{y}, a)}{\mathrm{Tech}_1(\mathrm{y}, a)}\left[1 - \frac{\mathrm{Dam}_1(d, y)}{\mathrm{AIS}(u, c)}\right] \quad (13)$$

As for direct production shocks in equation (9), prudence has to dominate the damage dependence of production (first term on right hand side). In addition, equation (13) 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

[19] Observe that except for Figure 12 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.

$\text{Dam}_1 > \text{AIS}(u, c)$ (reducing the right hand side of inequality 13). 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 (13) 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 (9) 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 factor productivity shocks (see previous section) to $\eta \in [0.6, 2]$ for mean-zero shocks on labor augmenting technology (assuming DICE's $\kappa = 0.3$).

A.4 Numeric Results for Analytic Discussion

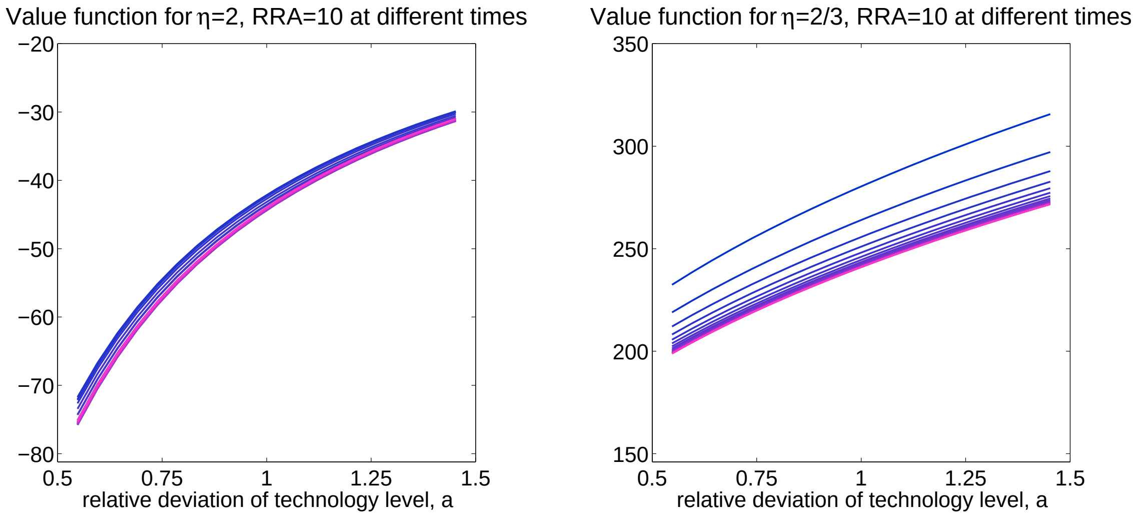

Figure 10 shows the normalized value function over the relative deviation of the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. We show two levels of consumption smoothing ($\eta = 2$, left panel and $\eta = 2/3$, right panel, both with RRA = 10). The 20 different lines correspond 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 $\eta = 2/3$.

Figure 10 presents the (normalized) expected welfare as a function of the technology level. It shows the normalized value function over the relative deviation of

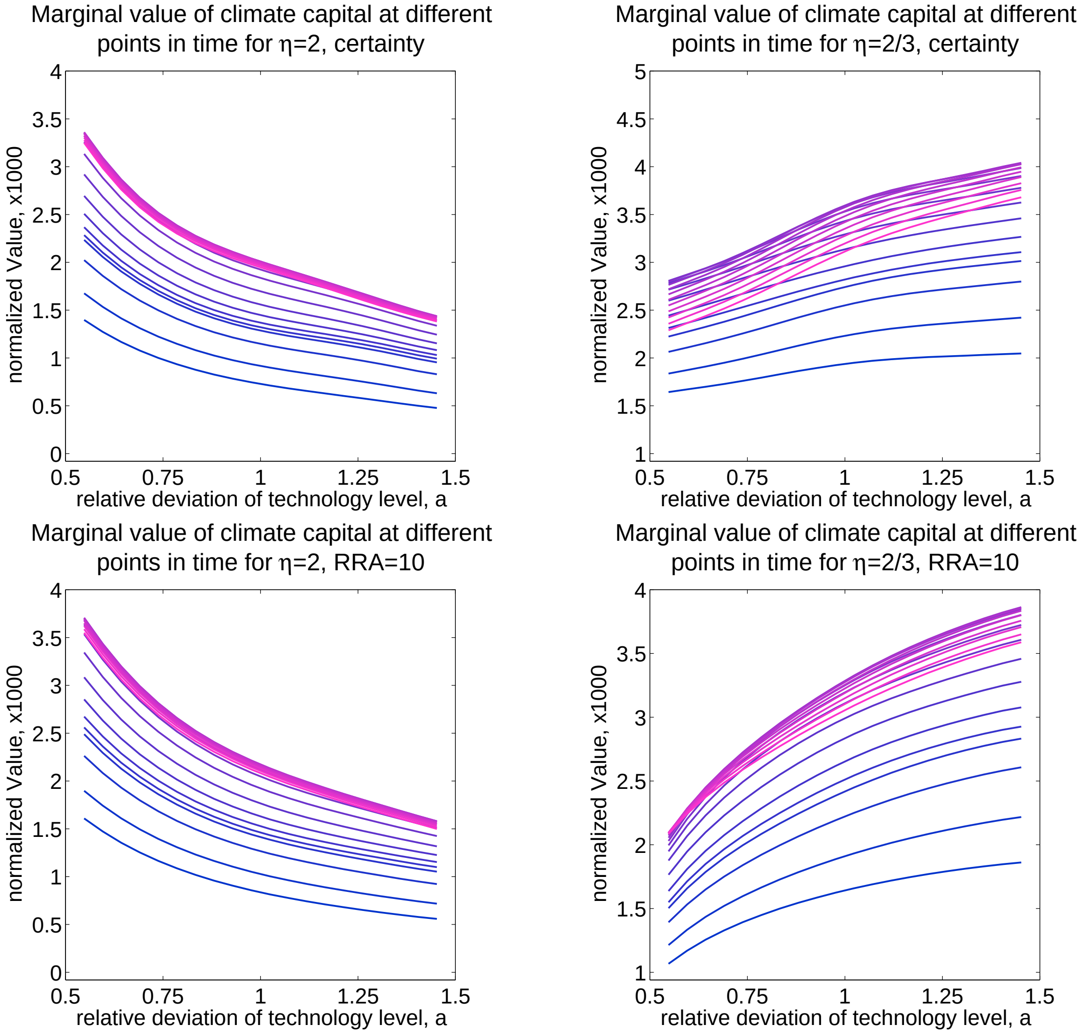

Figure 11 shows the marginal welfare gain $-\frac{\partial V_t^*}{\partial M_t}$ from an avoided ton of atmospheric carbon as a function of the normalized technology level $a = \frac{A}{A^{det}}$. In the case of a relatively high preference for consumption smoothing ($\eta = 2$, left panels) the marginal 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).

the technology level from its deterministic level, $a = \frac{A}{A^{det}}$. The two graphs correspond to the two disentangled preference scenarios with $\eta = 2, RRA = 10$ and with $\eta = 2/3, RRA = 10$. As in Figure 8, 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 $\eta = 2/3$.

Figure 11 shows 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 depicts 20 lines, each of which corresponds to a different point in time.

B Derivation of Analytic Results

This appendix derives the first order conditions and the analytic conditions characterizing the policy impact of a growth shock stated in sections 4 and A.3. We denote the constant consumption rate by α, and continue using the definitions of prudence $\text{Prud}(u, c)$, the aversion to intertemporal substitution $\text{AIS}(u, c)$, and the elasticities characterizing the marginal damage response to production $\text{Dam}_1(d, y)$ and $\text{Dam}_2(d, y)$ defined in sections 4.2 and 4.3, and the elasticities of (marginal) production with respect to technology $\text{Tech}_1(\text{y}, a)$ and $\text{Tech}_2(\text{y}, a)$ from section A.3.

Derivation of equation (10):
Optimizing the normalized Bellman equation (6) with respect to consumption returns

$$u'(c_t) = \beta_t \underbrace{\exp(-g_{A,t} - g_{L,t})}_{\equiv g_t} \underbrace{\frac{\mathbf{E}_t f'(V_{t+1})}{f'(f^{-1}\mathbf{E}_t f(V_{t+1}))}}_{\equiv \Pi_t} \; \mathbf{E}_t \underbrace{\frac{f'(V_{t+1})}{\mathbf{E}_t f'(V_{t+1})}}_{\equiv P_t} \frac{\partial V_{t+1}}{\partial k_{t+1}} \; . \tag{14}$$

In equation (10), the proportionality absorbs the growth factor g_t.

Derivation of equations (7) **and** (11):
Optimizing the normalized Bellman equation (6) with respect to abatement returns

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

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

$$\Rightarrow \Lambda'(\mu_t) = -\alpha(T_t, t) \; \beta_t \Pi_t \frac{\mathbf{E}_t \; P_t \; \frac{\partial V_{t+1}}{\partial M_{t+1}}}{u'(c_t)} \; . \tag{15}$$

In the last step we use equation (14). The second line corresponds to equation (11), except that in (11) the proportionality absorbs the first term on the right hand side, which only depends on the period t states. To derive equation (7), we continue and differentiate the Bellman equation (6) partially with respect to the carbon stock M_t

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

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

We applied the envelope theorem and simplified the expression using equation (14). To eliminate the value function from the right hand side, we repeatedly substitute this relation into itself, advancing the time indices period by period until we obtain the infinite sum stated below

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

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

$$= u'(c_t) \frac{\partial y_t}{\partial M_t} + \sum_{i=t}^{\infty} \left\{ \prod_{j=t}^{i} \beta_j \frac{\partial M_{j+1}}{\partial M_j} \Pi_j \mathbf{E}_j P_j \right\} \;\; u'(c_{i+1}) \frac{\partial y_{i+1}}{\partial M_{i+1}} \ . \tag{16}$$

We insert equation (16) into equation (15) resulting in

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

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

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

$$= -\frac{\alpha(T_t, t)}{\frac{\partial M_{t+1}}{\partial M_t}} \, \mathbf{E}_t^* \;\; \sum_{i=t}^{\infty} \left\{ \prod_{j=t}^{i} \beta_j \frac{\partial M_{j+1}}{\partial M_j} \Pi_j P_j \right\} \;\; \frac{u'(c_{i+1})}{u'(c_t)} \frac{\partial y_{i+1}}{\partial M_{i+1}} \ . \tag{17}$$

Equation (17) corresponds to equation (7). In equation (7), we used $\frac{\partial M_i}{\partial E_t} = \prod_{j=t}^{i} \frac{\partial M_{j+1}}{\partial M_j}$ and 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. Whereas the expectation operators $\mathbf{E}_t$ take expectations over the realization of $\tilde{A}_{t+1}$ conditional on earlier realizations of A_t, the operator $\mathbf{E}_t^*$ takes expectations over all possible future sequences $\tilde{A}_{t+1}, \tilde{A}_{t+2}, ...$ conditional on A_t.

Derivation of equation (8):
A deterministic technology shock in a given period increases the cost of carbon in

that period if $\frac{d}{da}u'(c(a))d(y(a))$ is positive.

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

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

Derivation of equation (9):
By Jensen's inequality, an uncertain mean-zero shock on production in a given period increases the cost of carbon in that period if $u'(\alpha y)d(y)$ is convex, i.e.,

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

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

$$\begin{aligned}-\frac{u'''(\alpha y)}{u''(\alpha y)}\alpha y = -\frac{u'''(c)}{u''(c)}c > 2\frac{d'(y)}{d(y)}y - \frac{\frac{d''(y)}{d(y)}y^2}{-\frac{u''(\alpha y)}{u'(\alpha y)}\alpha y} &= 2\frac{d'(y)}{d(y)}y - \frac{\frac{d''(y)}{d'(y)}y\frac{d'(y)}{d(y)}y}{-\frac{u''(c)}{u'(c)}c} \\ \mathrm{Prud}(u,c) \qquad > \qquad 2\ \mathrm{Dam}_1(d,y)\ &- \frac{\mathrm{Dam}_1(d,y)}{\mathrm{AIS}(u,c)}\ \mathrm{Dam}_2(d,y)\ ,\end{aligned}$$

which is equation (9).

Derivation of equation (13) in Appendix A.3:
By Jensen's inequality, an uncertainty mean-zero shock on the rate of technological growth in a given period increases the cost of carbon contribution in that period if the product $u'(\alpha y(a))d(y(a))$ is convex, i.e.,

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

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

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

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

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

This last equation coincides with equation (13) in Appendix A.3.

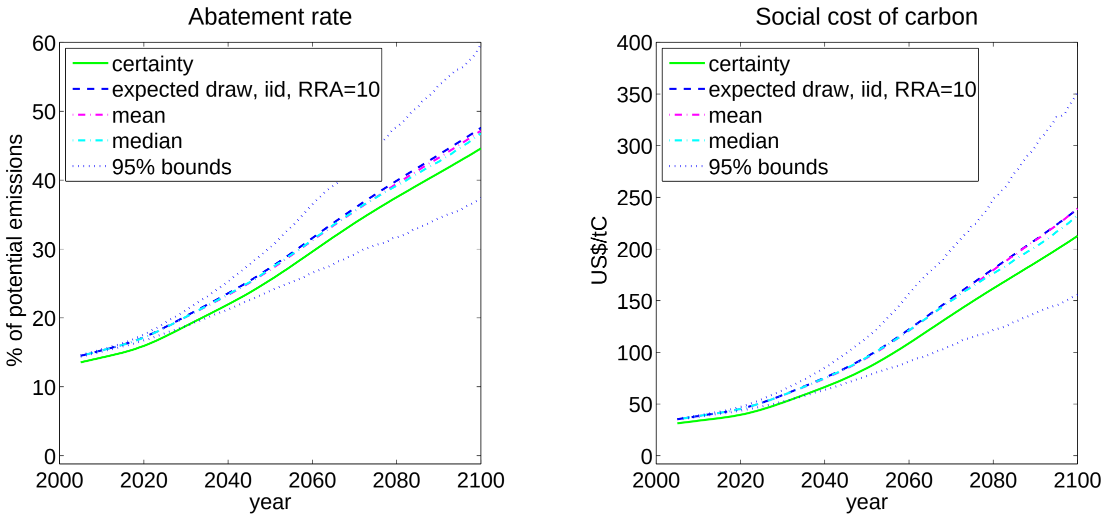

Figure 12 shows the mean, the median, the expected draw representation, and the 95 % confidence bounds of 1000 optimal random paths for abatement and the social cost of carbon. The social preference parameters are $\eta = 2, RRA = 10$. We compare the different measures to the the optimal paths under 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.

Online Appendix

I Further Results

Our figures in the main text represent uncertainty by expected draws. Figure 12 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.

In DICE, for the base specification where $\eta = 2$, a deterministic increase in the technology level reduces the optimal carbon tax. Figure 13 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 13 shows the optimal abatement rate and the right panel shows the optimal social cost of carbon. 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 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 (8) also holds in the future, it assumes a given state of the climate system. Comparing the different lines in Figure 13, the carbon stock during the second half of the century is significantly

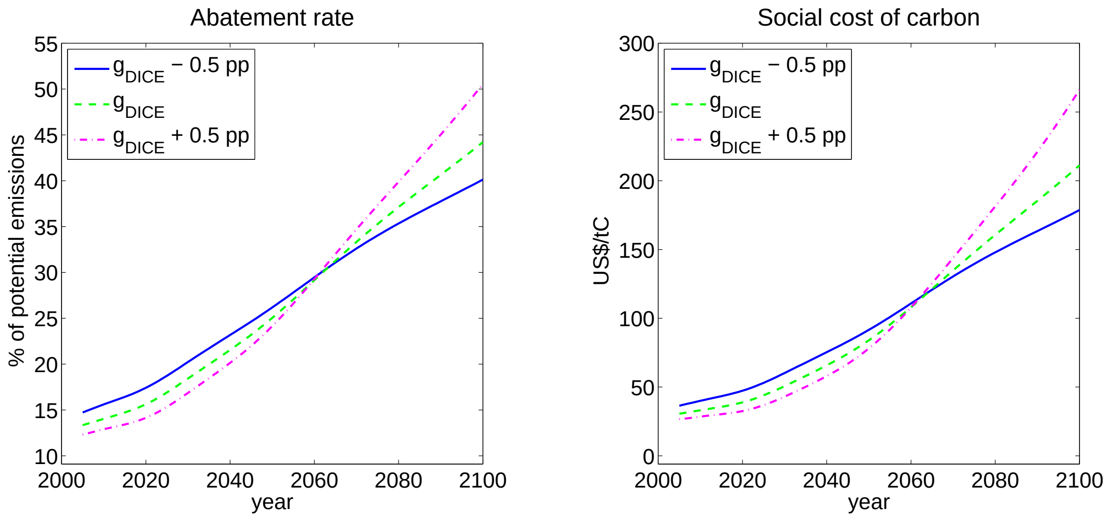

Figure 13 compares the optimal abatement rate and social cost of carbon under certainty for three growth rates: the DICE growth rate, the DICE growth rate +0.5 percentage points, and the DICE growth rate −0.5 percentage points. Higher growth lowers the optimal present-day abatement but increases abatement steeply later in the century.

higher for the generations on the high growth path as compared to those on the low growth path.

Figures 14 and 15 show the result of calibrating our simplified climate module to the original DICE-2007 model. Figure 14 shows the case of standard preferences ($\eta = 2$), whereas in Figure 15 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.

II Renormalizing the Bellman Equation and Numeric 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:[20]

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

This transformation also concentrates the Chebychev nodes at which we evaluate our Chebychev polynomials in the close future in real time, where most of the exogenously driven changes take place. Further, we improve the performance of the recursive numeric 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, $z_t = 0 \ \forall t$)

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

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

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

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

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

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

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

which leads to the new Bellman equation (6)

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

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 numeric 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

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

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

We recover the original value function from

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

The marginal value of a ton of carbon is given by

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

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

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

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

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

III The Climate Enriched Economy Model

The following model is largely a reproduction of DICE-2007 (Nordhaus 2008) and its recursive dynamic programming implementation in Traeger (2012). The three most notable differences with respect to DICE-2007 are the annual time step (DICE-2007 features a decadal step), the infinite time horizon, and a simplification of the carbon cycle and temperature delay structure in our model. Traeger (2012) discusses the simplification of the carbon cycle in detail and tests it against the MAGICC 6.0 model, which emulates the big Atmosphere-Ocean General Circulation Models used by the Intergovernmental Panel on Climate Change. In contrast to Traeger (2012), we further simplify the temperature delay structure to save an additional state variable. State variables are computationally very expensive in a dynamic programming implementation because of the curse of dimensionality, and we need two additional states for technology and shock persistence as compared to the baseline model in Traeger (2012). All parameters are characterized and quantified in Table B 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 14 and 15) in Online Appendix I.

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

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

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

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

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 14 and 15 compare the performance of our model to the original DICE model.

IV Growth Rate Shocks

This section discusses our implementation of growth shocks. First, we discuss general properties of the implementations. Second, we show that our iid shock implementation delivers, in expectation, the same growth trajectory as our deterministic base case. Third, we show this coincidence of expected technology level and deterministic technology evolution for the case of persistent shocks.

IV.1 Background

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 the current technology level at time t by A_t and the uncertain future technology level by $\tilde{A}_{t+1}$. Let A^{det}_{t+1} denote the growth trend, i.e., the (hypothetical) deterministic technology level in the absence of growth shocks in all periods. Finally, $a_t = A_t/A^{det}_t$ measures the deviation of the actual technology level from the deterministic level. 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 the growth shock. 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 $t+j$ is cumulative in the growth shocks:

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

IV.2 Technology expectation under iid shocks

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

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

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

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

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

IV.3 Technology expectation under 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 $\tilde{y}_{t+j}$ can be written as

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

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

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

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

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

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

Equation (21) tells us how the 'cumulative' variance of the normal distribution in the exponent for $\tilde{A}_{t+j}$ increases over time. The factors ${}^{1-\gamma^{j-i+1}}/_{1-\gamma} > 1$ increase in j 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.

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

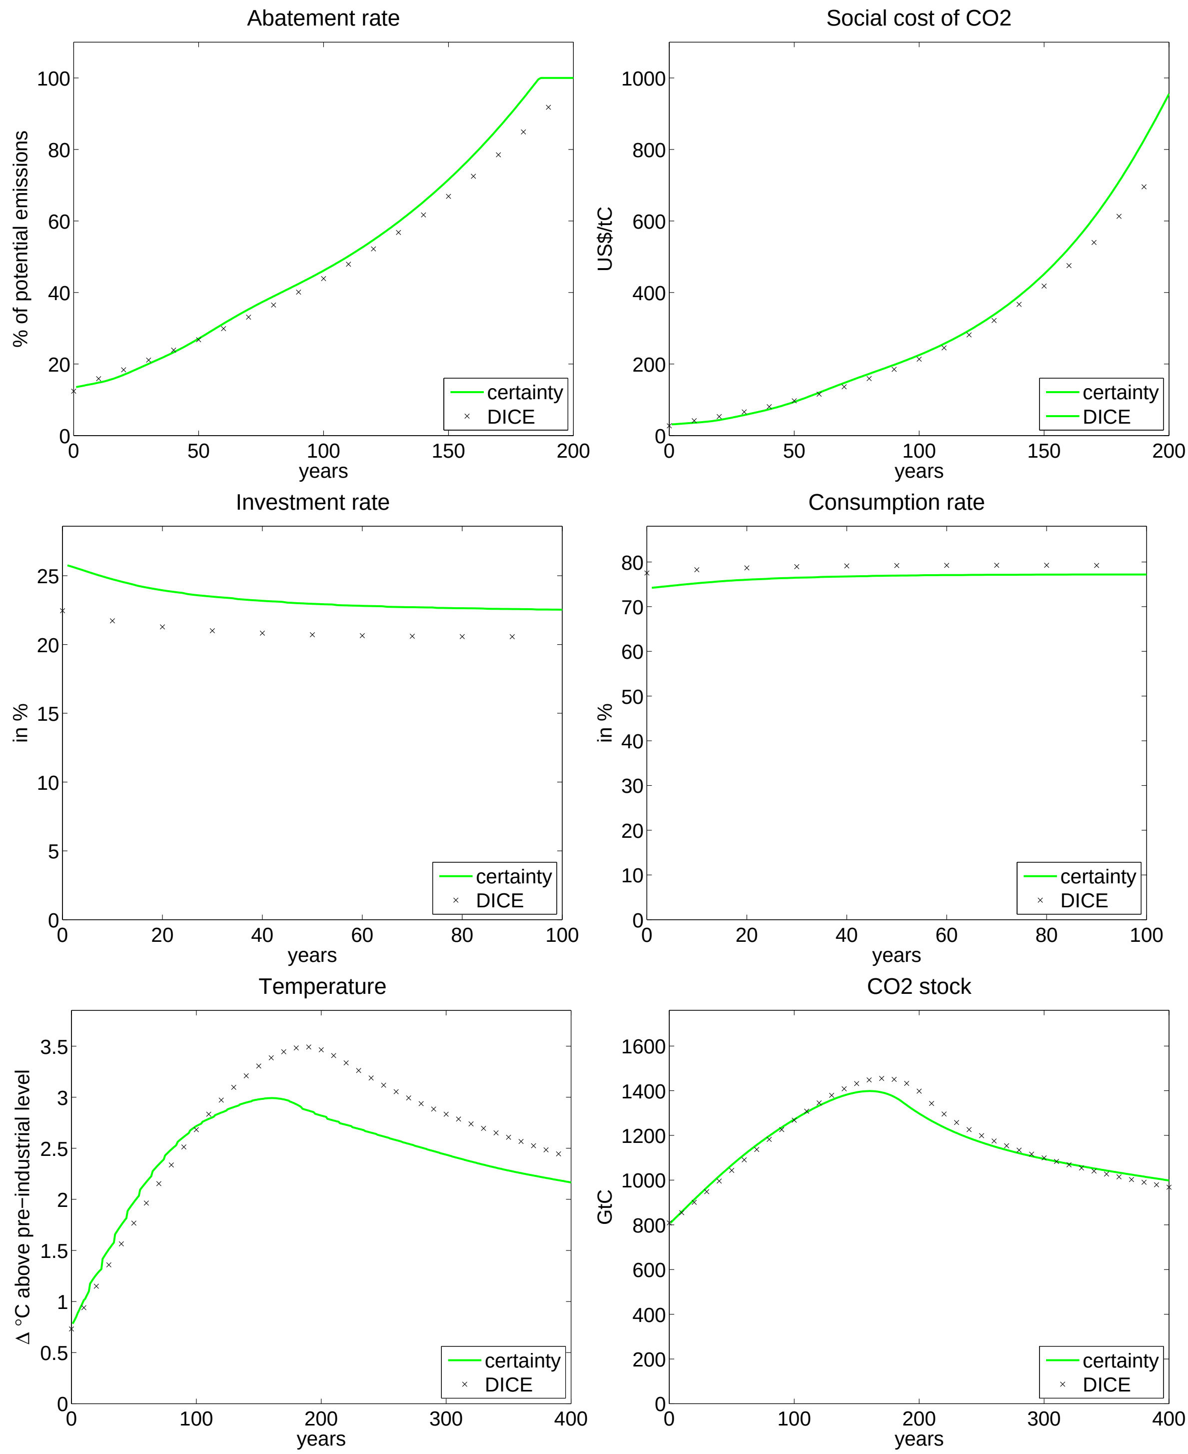

Figure 14 compares the results our recursive formulation under certainty with the original DICE model results The consumption smoothing parameter is $\eta = 2$.

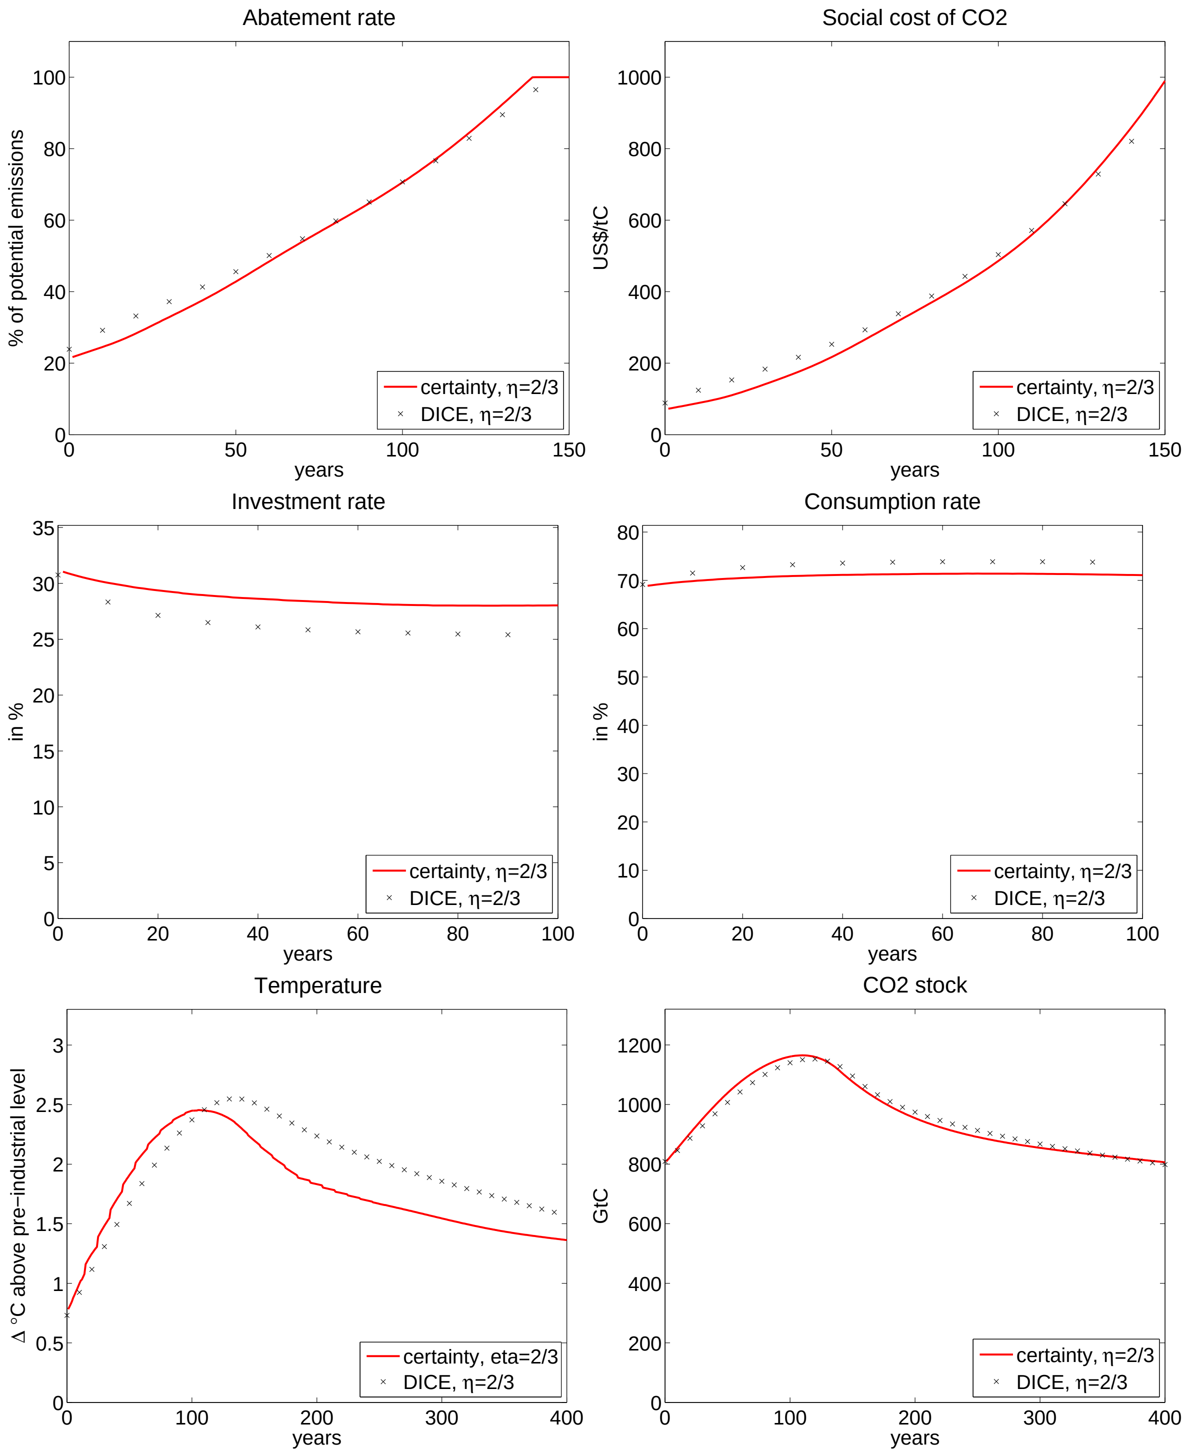

Figure 15 compares the results of our recursive formulation under certainty with the DICE model results for the same low consumption smoothing parameter $\eta = 2/3$.