

ACE – Analytic Climate Economy (with Temperature and Uncertainty)

Christian P. Traeger

Department of Economics, University of Oslo

December 2017

Preliminary revision prepared for ASSA 2018

– Please cite SSRN version to be released/updated soon –

Abstract: The paper develops an analytic integrated assessment model of climate change. It enhances our current understanding of climate policy and explains crucial relations to the broader audience. The model offers a novel framework to address climate change uncertainties. The analytic solution overcomes Bellman’s curse of dimensionality for a wide range of stochastic processes. I analyze the policy implications of the main climate uncertainties and show the different welfare implications of “objective” uncertainty, epistemological uncertainty, and anticipated learning. In contrast to earlier suggestions in the literature, uncertainty is not more relevant to climate change evaluation than discounting, but uncertainty makes the policy recommendations even more sensitive to the calibration of the discount rates and its individual components than under certainty. The present Analytic Climate Economy (ACE) is the first analytic model comprising all the components considered essential for quantitative policy advising.

JEL Codes: Q54, H43, E13, D81, D90, D61

Keywords: climate change, integrated assessment, uncertainty, learning, risk aversion, recursive utility, social cost of carbon, carbon tax, carbon cycle, climate sensitivity, stochastic volatility, autoregressive gamma

1 Introduction

Integrated assessment of climate change analyzes the interactions of long-term economic growth, greenhouse gas emissions, and global warming. The present analytic climate economy (ACE) is a quantitative model competing with numeric models used to derive the US federal social cost of carbon. The analytic solution permits new insights into the evaluation of climate change, and it overcomes numeric obstacles in incorporating uncertainties. The reader can evaluate the implications of changing contested parameters by immediate inspection or changing simple formulas on a spreadsheet. The analytic nature of the present quantitative model attempts to reconnect the integrated assessment community's contributions published in field journals, the abstract climate thoughts published in general journals, and the broad economic audience. An increasing number of economists is interested in the biggest environmental problem of our generation, but hesitant to penetrate a world of complex numeric models that appear to be black boxes to the outsider.

The paper's focal point is the welfare and policy response to deterministic and uncertain climate dynamics. It relates the optimal carbon tax directly to the characteristics of the carbon cycle and the climate system. Recent stylized models find that uncertainty surrounding climate change could be the major driver of welfare loss and mitigation policy. Policy advising remains in the hands of deterministic models that explore and average large samples of deterministic worlds. In these Monte-Carlo runs the decision maker is not aware of the uncertainty and sees only one world at a time. In contrast, stochastic models integrate uncertainty into the model, the welfare, and the decision making. Numeric stochastic integrated assessment models (IAMs) are on the rise and have delivered major insights over the recent years. New methods lessen Bellman's curse of dimensionality, yet the dimensionality of integrated assessment model still seriously restricts numeric analysis.¹ These numeric limitations make an analytic model like ACE particularly valuable as it can incorporate more of the relevant climatic and informational states than contemporaneous stochastic numeric IAMs.

The most frequently discussed uncertainty in climate change economics is the climate's sensitivity to atmospheric carbon dioxide concentrations. We currently do not know whether a doubling of atmospheric carbon dioxide from its pre-industrial level will lead to a 1.5°C or a 4.5°C warming (IPCC's likely range) or even more. We have about 70% of the greenhouse gases for an equivalent doubling already in the atmosphere, yet climate inertia is delaying their impacts. Climate science itself pays even more attention to the uncertainty governing atmospheric carbon build-up. Carbon dioxide does not decay, it only moves between

¹These numeric stochastic dynamic programming models solve the full non-linear system on a state space containing both economic and climate variables. A linearization around a steady state is generally unsatisfactory in a transiting climate economy with many non-linear interactions. The present model's analytic solution method also suggests a promising approach for reducing the numeric "curse" in related models that defy a closed-form solution.

reservoirs including the atmosphere, the oceans, and the biosphere (plants and soils). The precise carbon flows between these reservoirs are largely uncertain. For example, an amount of carbon weighing more than the entire human race walking this planet vanishes every year from the atmospheric carbon budget into an unidentified sink. We do not know whether this carbon will continue to leave the atmosphere in a warming climate, whether its flow will stall, or whether it might return back into the atmosphere.

Economic models convey ideas and insights, here, governing the the main drivers of climate policy and their interaction. Discount rates have long been assumed to be the main determinant of optimal climate policy, changing optimal climate policy targets much more than improved estimates of emissions, damages, or warming. Recently, Pindyck (2013) argued that uncertainty is *the* crucial characteristic of climate change that outweighs all other components and Weitzman (2009*b*) argues that fat-tailed temperature uncertainty makes the choice of time preference largely irrelevant, an argument countered by Roe & Bauman (2013) because of the major delay in climate response. I show that the welfare loss from uncertainty is even more sensitive to discounting than its deterministic contribution, i.e., the contribution from expected change. The sensitivity to nature's uncertainty increases with the power of the distributional moments. Thus, fat tails increase the sensitivity to discount rates instead of reducing it.

Uncertainties in climate change have always called the climate skeptics on the table, promoting a wait and see policy. Various studies demonstrate that learning is too slow to substantially affect the optimal carbon policy (Kelly & Kolstad 1999, Leach 2007, Jensen & Traeger 2013, Gerlagh & Liski 2014, Kelly & Tan 2015*a*). The present paper derives analytic insights into the implications of a stochastic climate, of epistemological uncertainty (scientific lack of knowledge), and of anticipated learning. I show that nature's stochasticity and epistemological uncertainty imply opposing sensitivities to time preference. Moreover, knowledge updates make a Bayesian learning framework the most sensitive to time preference because updates change the long-run picture of the future.

Guided by the long run risk literature, ACE disentangles risk aversion from consumption smoothing to calibrate the risk-free discount rate and risk premia separately. Models lacking this feature are forced to either discount the future too highly, or to disrespect the risk premia. I show that the relevant risk aversion for climate change evaluation is not Arrow Pratt's measure of risk aversion, but by how much Arrow Pratt risk aversion exceeds the desire to smooth consumption over time (intrinsic aversion to risk). Higher moments of the uncertainty distribution are evaluated with higher powers of such risk aversion.

Economic models guide discussions and help in quantifying policy targets. The discussion of the discount rate in the climate context is prominent for a good reason. It remains the most relevant determinant of the optimal carbon tax. It also remains the input over which economists are most divided. The analytic solution makes it easy to adjust parameters to reflect individual perspectives, and to translate philosophical differences or different calibra-

tion approaches into their quantitative policy implications. Ethical arguments as well as the long run risk model’s calibration to asset prices lead to a low rate of pure time preference. ACE’s analytic nature permits an easy solution for low rates of pure time preference. In deterministic models, these low discount rates require very long and computationally expensive time horizons beyond the common model specifications. In stochastic models, a low discount rate reduces the contraction of the Bellman equation and numeric issues frequently prevent a solution. The model also shows why pure time preference can be much more important for long-term climate change evaluation than the consumption discount rate.

ACE is the first IAM that splits up the carbon tax and welfare loss contributions between the carbon cycle’s and the climate system’s contributions. Under certainty, the carbon cycle is the main driver. The persistence of atmospheric carbon increases the optimal carbon tax by a factor of 3-30 (depending on pure time preference), as compared to a 15-40% reduction resulting from a combination of warming delay and temperature persistence. Under uncertainty, the relevance of the two components flips. Due to the non-linearities in the interactions, carbon flow uncertainty plays a minor role whereas temperature uncertainty raises the welfare loss and carbon tax substantially. These findings have immediate implications for research priorities in the climate sciences.

Economic models guide research. Agreement on a single integrated assessment model is not on the horizon, and likely not desirable either. ACE is a valuable contribution to the set of models as it fleshes out assumptions of standard models and their implications. Some of these assumptions will be, should be, and have been challenged. Building and discussing new insights is easier with the transparency of an analytically tractable model. It can accompany numeric extensions with approximate interpretations of the changes. The analytic benchmark serves as a tool for insight, quick quantification, and as a platform to challenge ideas; and it can play the scapegoat, transparently representing common assumptions and, thereby, furthering new research directions.

1.1 ACE’s Relation to Other Analytic IAMs

Analytic approaches to the integrated assessment of climate change go back at least to Heal’s (1984) insightful non-quantitative contribution. A series of papers has used the linear quadratic model for a quantitative analytic discussion of climate policy (Hoel & Karp 2002, Karp & Zhang 2006, Karp & Zhang 2012). In linear quadratic models *welfare* responds to uncertainty. In the wide-spread additive noise model, optimal *policy* remains unaffected by risk (weak certainty equivalence). In Hoel & Karp’s (2001) multiplicative noise model also the optimal policy responds to uncertainty. A disadvantage of the linear quadratic model is its highly stylized representation of the economy and the climate system. In particular, the model has no production or energy sector. Recently, Golosov et al. (2014) broke new ground by amending the log-utility and full depreciation version of Brock & Mirman’s (1972)

stochastic growth model with an energy sector and an impulse response of production to emissions.

Golosov et al.’s (2014) elegant model makes use of two climate change characteristics. First, a decadal time step is neither uncommon in IAMs nor particularly problematic given the time scales of the climate change problem. Therefore, the full-depreciation assumption is much more reasonable than in other macroeconomic contexts. The present paper further weakens the full depreciation assumption. Second, planetary “heating” (radiative forcing) is logarithmic in atmospheric carbon and damages are convex in temperature. As a consequence, the authors argue for a linear relation between past emissions and present damages. Their argument assumes that temperature responds immediately to atmospheric carbon increase. However, reaching a new equilibrium temperature after increases in atmospheric carbon can take decades to centuries. Gerlagh & Liski (2012) extend the model by introducing the empirically important delay between emission accumulation and damages. The present paper follows the numeric IAMs used in policy advising and explicitly introduces the logarithmic relation between carbon dioxide’s radiative forcing and temperature change (Nordhaus 2008, Hope 2006, Anthoff & Tol 2014). Moreover, I incorporate a novel model of ocean-atmosphere temperature dynamics that competes well with these numeric policy models. As an additional payoff, ACE is the first analytic IAM to define and calibrate damages on temperature rather than on carbon. Integrating the non-linear relations between carbon and temperature is crucial for an application to climate uncertainty.

Golosov et al.’s (2014) framework has been used to examine a multi-regional setting (Hassler & Krusell 2012), non-constant discounting (Gerlagh & Liski 2012, Iverson 2013), intergenerational games (Karp 2013), and regime shifts Gerlagh & Liski (2014). Golosov et al.’s (2014) framework imposes *strong certainty equivalence*: not even welfare responds to uncertainty. I show that this feature arises from simultaneously setting the intertemporal elasticity of substitution and Arrow Pratt risk aversion to unity. Whereas unity is within the estimated range of the intertemporal elasticity of substitution, Arrow Pratt risk aversion is ubiquitously estimated higher. I solve ACE for arbitrary degrees of (disentangled) Arrow Pratt risk aversion, accommodating for one of the most prominent criticisms of the model. Constant relative Arrow Pratt risk aversion implies a decreasing coefficient of absolute risk aversion. This stylized fact is widely believed to hold and contrasts with linear quadratic AIAMs that only capture increasing absolute Arrow Pratt risk aversion or risk neutrality. Alternatively, Li et al. (2014) and Anderson et al. (2014) leave the world of von Neumann & Morgenstern’s (1944) axioms and introduce a preference for robustness to escape the strong certainty equivalence of the Golosov et al. (2014) framework.² The present paper breaks with both strong and weak certainty equivalence, while maintaining von Neumann &

²Anderson et al. (2014) deviate from Golosov et al. (2014) by using a linear relation between the economic growth rate, temperature increase, and cumulative historic emissions. Both Li et al. (2014) and Anderson et al. (2014) combine a simpler analytic model with a more complex numeric IAM for quantitative simulation.

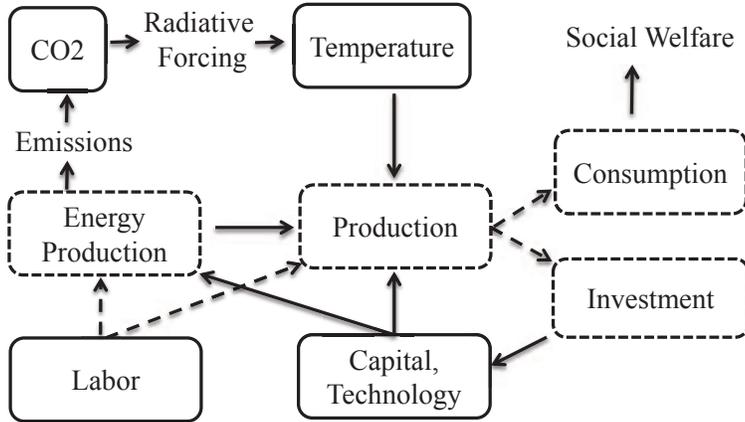


Figure 1: The structure of ACE and most Integrated Assessment Models. Solid boxes characterize the model’s state variables, dashed boxes are flows, and dashed arrows mark choice variables.

Morgenstern’s (1944) classic axioms for choice under uncertainty, which are often considered desirable for rational or normative choice.

2 The Model

ACE’s structure follows that of most IAMs (Figure 1). Labor, capital, and technology create production that is either consumed or invested. Production relies on energy inputs which cause emissions. Emissions accumulate in the atmosphere, cause radiative forcing (greenhouse effect), and increase global temperature(s), reducing production. This section introduces the basic model of the economy, the energy sector, and the climate system. It derives the necessary and sufficient assumptions to solve the model in closed form and summarizes the underlying calibration.

2.1 The Basic ACE Economy

Utility is logarithmic and the social planner’s time horizon is infinite. The logarithmic utility function captures only consumption smoothing over time. I assume a stable population normalized to unity, but the approach generalizes to a population weighted sum of logarithmic per capita consumption with population growth. Gross production is a Cobb-Douglas function of technology level $A_{0,t}$, capital K_t , the energy composite E_t , and the amount of labor $N_{0,t}$ employed in the final consumption good sector

$$Y_t^{net} = A_{0,t} K_t^\kappa N_{0,t}^{1-\kappa-\nu} E_t^\nu . \quad (1)$$

The aggregate energy input E_t is a smooth and monotonic function

$$E_t = g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) \quad (2)$$

of $I \in \mathbb{N}$ different energy sources, whose production levels $E_{i,t}$ are collected in the vector $\mathbf{E}_t \in \mathbb{R}_+^I$. These decomposed energy inputs are produced using technologies $\mathbf{A}_t \in \mathbb{R}_+^I$ and labor input levels $\mathbf{N}_t \in \mathbb{R}_+^I$. Total labor supply is normalized to unity, $\sum_{i=0}^I N_{i,t} = 1$. The first I^d energy sources are fossil fuel based and emit CO₂ (“dirty”). I measure these energy sources in units of their carbon content. Their extraction is costly, they are potentially scarce, and I denote this subset of energy inputs by the vector $\mathbf{E}_t^d \in \mathbb{R}_+^{I^d}$. Total emissions from production amount to $\sum_{i=1}^{I^d} E_{i,t}$. Renewable energy sources indexed $I^d + 1$ to I are costly but not scarce, and their production does not emit CO₂ (“clean”). I assume a system of energy sectors of the general form (2) that is sufficiently smooth and well-behaved to let the value function converge and to avoid boundary solutions.³

The dirty fossil fuel energy sources are (potentially) scarce and their resource stock in the ground $\mathbf{R}_t^d \in \mathbb{R}_+^{I^d}$ follows the equation of motion

$$\mathbf{R}_{t+1} = \mathbf{R}_t - \mathbf{E}_t^d,$$

with initial stock levels $\mathbf{R}_0 \in \mathbb{R}_+^{I^d}$ and $\mathbf{R}_t \geq 0$ at all times. The next section explains how the energy sector’s carbon emissions increase the global atmospheric temperature $T_{1,t}$ measured as the increase over the preindustrial temperature level. This temperature increase causes damages, which destroy a fraction $D_t(T_{1,t})$ of production. I normalize $D_t(0) = 0$ and Proposition 1 characterizes the class of damage functions $D_t(T_{1,t})$ that permit an analytic solution of the model.

Weakening Golosov et al.’s (2014) assumption of full depreciation, I assume the capital stock’s equation of motion

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})](1 - x_t) \left[\frac{1 + g_{k,t}}{\delta_k + g_{k,t}} \right], \quad (3)$$

where $x_t = \frac{C_t}{Y_t[1 - D_t(T_{1,t})]}$ is the endogenous consumption rate and $g_{k,t}$ is an exogenous approximation of the growth rate of capital. The consumption rate replaces absolute consumption as the consumption-investment control to achieve additive separability between controls and states for log capital. If either $g_{k,t} = \frac{K_{t+1}}{K_t} - 1$ (actual capital growth rate) or $\delta_k = 1$ (full depreciation), then equation (3) coincides with the standard assumption on capital accumulation (see Appendix B)

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})] - C_t + (1 - \delta_k)K_t .$$

The depreciation correction $\frac{1+g_{k,t}}{\delta_k+g_{k,t}}$ is larger the slower the capital stock depreciates and the slower it grows. The factor makes the decision maker aware of the additional capital

³Sufficient but not necessary conditions are smoothness, interior solutions for the controls, and convexity of the energy production set. See Golosov et al. (2014) for a three energy sector example, and note that in the present setting substitutabilities can change over time accounting for progress in electricity storage improving the substitutability between liquid fuels, coal, and renewables.

in the next period and it can adjust ACE’s capital-output ratio and capital depreciation to macroeconomic observation, addressing a critique raised against Golosov et al.’s (2014) model.⁴

Yet, the factor has no impact on the optimal carbon policy, given current world output. The relevant implication of the capital accumulation in equation (3) and the full depreciation assumption is that the investment rate is independent of the system states. As a consequence, climate policy will not operate through the consumption rate. Appendix B shows that the consumption rate is approximately independent of the climate states also in an annual time step version of the DICE model (and as well for non-logarithmic utility).

2.2 The Extended ACE Economy: Capital Persistence, Endogenous Growth, and Capital Damages

I present a simple extension of the “base ACE” that incorporates capital persistence and simple endogenous growth. It avoids the exogenous approximation or full depreciation assumptions above. The model introduces a dedicated capital production sector

$$V_t = A_{-1,t} K_{-1,t}^{\bar{\kappa}} N_{-1,t}^{1-\bar{\kappa}-\bar{\nu}} E_{-1,t}^{\bar{\nu}}.$$

The investment good V_t is linear homogenous in capital, labor, and energy input in line with the replication argument. The subindex -1 on these inputs labels the level dedicated to capital production. A standard AK model reasoning would close the capital sector with the assumption that innovations increase the freely available technology level $A_{-1,t}$ (ideas, blue prints,...) as $A_{-1,t} = A_{-1,0} K_t^{1-\bar{\kappa}}$. In ACE’s climate economy the energy sector can play a more crucial role than in standard AK models, generalizing the classical ? assumption to

$$A_{-1,t} E_{-1,t}^{\bar{\nu}} = A_{-1,0} K_t^{1-\bar{\kappa}} \bar{E}_{-1,t}^{\bar{\nu}},$$

where $\bar{E}_{-1,t}$ labels the detrended energy composite with long-term growth removed. This detrended component $\bar{E}_{-1,t}$ can respond to energy shocks and energy transition, and it can affect the growth rate of the underlying economy.

The capital’s equation of motion, replacing equation (3) of the “base ACE”, is

$$K_{t+1} = [(1 - \delta)K_t + V_t][1 - D_K(T_{1,t})],$$

introducing a dedicated damage function $D_K(T_{1,t})$ for the capital sector. Equation (1) now characterizes production of an aggregate consumption commodity, and the capital input into equation (1) changes to $K_t - K_{-1,t}$ because capital input is split between final good and capital production.

⁴I note that already the model’s time step of 10 years makes the capital depreciation assumption more reasonable than it might appear: instead of an annual decay that leaves 30%-40% after 10 years, the model utilizes all of the capital during 10 years, and none afterwards.

2.3 The ACE Climate System

The energy sector’s CO₂ emissions enter the atmosphere. Land conversion, forestry, and agriculture also emit smaller quantities of CO₂. Following the DICE model, I treat these additional anthropogenic emission as exogenous and denote them by E_t^{exo} . Carbon released into the atmosphere does not decay, it only cycles through different carbon reservoirs. Let $M_{1,t}$ denote the atmospheric carbon content and $M_{2,t}, \dots, M_{m,t}$, $m \in \mathbb{N}$, the carbon content of a finite number of non-atmospheric carbon reservoirs that exchange carbon. DICE uses two carbon reservoirs besides the atmosphere: $M_{2,t}$ captures the combined carbon content of the upper ocean and the biosphere (mostly plants and soil) and $M_{3,t}$ captures the carbon content of the deep ocean. Scientific climate models often use additional reservoirs. The vector \mathbf{M}_t comprises the carbon content of the different reservoirs and the matrix Φ captures the transfer coefficients. Then

$$\mathbf{M}_{t+1} = \Phi \mathbf{M}_t + \mathbf{e}_1 (\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo}) \quad (4)$$

captures the carbon dynamics. The first unit vector \mathbf{e}_1 channels new emissions from fossil fuel burning $\sum_{i=1}^{I^d} E_{i,t}$ and from land use change, forestry, and agriculture E_t^{exo} into the atmosphere $M_{1,t}$.

An increase in atmospheric carbon causes a change in our planet’s energy balance. In equilibrium, the planet radiates the same amount of energy out into space that it receives from the sun. Atmospheric carbon $M_{1,t}$ and other greenhouse gases (GHGs) “trap” some of this outgoing infrared radiation, which causes the (additional, anthropogenic) radiative forcing

$$F_t = \eta \frac{\log \frac{M_{1,t} + G_t}{M_{pre}}}{\log 2} . \quad (5)$$

The exogenous process G_t captures non-CO₂ greenhouse gas forcing measured in CO₂ equivalents. There is no anthropogenic radiative forcing if $G_t = 0$ and $M_{1,t}$ equals the preindustrial atmospheric CO₂ concentration M_{pre} . We can think of radiative forcing as a small flame turned on (or up) to heat a big pot of soup (our planet with its oceans). The parameter η captures the strength of this flame for a doubling of CO₂ with respect to the preindustrial concentration M_{pre} . Whereas radiative forcing is immediate, the planet’s temperature (the big pot of soup) reacts with major delay. After several centuries, the new equilibrium⁵ temperature caused by a new level of radiative forcing F^{new} will be $T_{1,eq}^{new} = \frac{s}{\eta} F^{new} = \frac{s}{\log 2} \log \frac{M_{1,eq} + G_{eq}}{M_{pre}}$. The parameter s is known as climate sensitivity. It measures the medium to long-term temperature response to a doubling of preindustrial CO₂ concentrations. Its best estimates lie

⁵The conventional climate equilibrium incorporates feedback processes that take several centuries, but excludes feedback processes that operate at even longer time scales, e.g., the full response of the ice sheets.

currently around 3C, but the true temperature response to a doubling of CO₂ is highly uncertain.

Next period's atmospheric temperature depends on the current atmospheric temperature, the current temperature in the upper ocean, and on radiative forcing. Next period's temperature in the upper ocean depends on current temperature in the adjacent layers: the atmosphere and the next lower ocean layer. I denote the temperature of a finite number of ocean layers by $T_{i,t}, i \in \{2, \dots, l\}, l \in \mathbb{N}$. I abbreviate the atmospheric equilibrium temperature resulting from the radiative forcing level F_t by $T_{0,t} = \frac{\varepsilon}{\eta} F_t$. A given ocean layer slowly adjusts its own temperature to the temperature of the surrounding layers. I model next period's temperature in layer $i \in \{1, \dots, l\}$ as a generalized mean of its present temperature $T_{i,t}$ and the present temperatures in the adjacent layers $T_{i-1,t}$ and $T_{i+1,t}$ ⁶

$$T_{i,t+1} = \mathfrak{M}_i^\sigma(T_{i,t}, w_i^{-1}T_{i-1,t}, w_{i+1}T_{i+1,t}) \text{ for } i \in \{1, \dots, l\}. \quad (6)$$

The weight matrix σ characterizes the (generalized) heat flow between adjacent layers, and σ^{forc} characterizes the heat influx response to radiative forcing. The equilibrium temperature ratios $w_i = \frac{T_{i-1,eq}}{T_{i,eq}}$ are empirical adjustments reflecting that the equilibrium warming does not coincide across all layers: in a warmer equilibrium the oceans lose more energy through evaporation, keeping them cooler relative to the atmosphere. Based on the data, my calibration in section 2.5 adjusts only for the equilibrium warming difference between atmosphere and oceans ($w_i = 1$ for $i \neq 2$). Proposition 1 in the next section characterizes the class of means (weighting functions f) that permit an analytic solution.

2.4 Solving ACE

Appendix C.3 solves ACE by transforming it into an equivalent linear-in-state model (Karp 2013). This transformation helps to understand which extensions maintain (or destroy) its analytic tractability. Linear-in-state models rely on equations of motion that are linear in the state variable, and on control variables that are additively separable from the states. ACE is linear only after transforming some of the original state variables. The policymaker optimizes labor inputs, consumption, and investment to maximize the infinite stream of logarithmic utility from consumption, discounted at factor β , over the infinite time horizon. The present paper assumes that the optimal labor allocation has an interior solution and

⁶A generalized mean is an arithmetic mean enriched by a non-linear weighting function f . It takes the form $\mathfrak{M}_i(T_{i-1,t}, T_{i,t}, T_{i+1,t}) = f^{-1}[\sigma_{i,i-1}f(T_{i-1,t}) + \sigma_{i,i}f(T_{i,t}) + \sigma_{i,i+1}f(T_{i+1,t})]$ with weight $\sigma_{i,i} = 1 - \sigma_{i,i-1} - \sigma_{i,i+1} > 0$. The weight $\sigma_{i,j}$ characterizes the (generalized) heat flow coefficient from layer j to layer i . Heat flow between any two non-adjacent layers is zero. Note that the weight $\sigma_{i,i}$ captures the warming persistence (or inertia) in ocean layer i . The weight $\sigma_{1,0} = \sigma^{forc}$ determines the heat influx caused by radiative forcing. I define $\sigma_{l,l+1} = 0$: the lowest ocean layer exchanges heat only with the next upper layer. For notational convenience equation (6) below writes a mean of three temperature values also for the deepest layer ($i = l$) with a zero weight on the arbitrary entry $T_{l+1,t}$. I collect all weights in the $l \times l$ matrix σ , which characterizes the heat exchange between the atmosphere and the different ocean layers.

that scarce resources are stretched over the infinite time horizon along the optimal path, avoiding boundary value complications. Linear-in-state models are solved by an affine value function. The following proposition summarizes the main result of Appendix C.3.

Proposition 1 *An affine value function of the form*

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t$$

solves ACE if, and only if, $k_t = \log K_t$, $\boldsymbol{\tau}_t$ is a vector composed of the generalized temperatures $\tau_{i,t} = \exp(\xi_i T_{i,t})$, $i \in \{1, \dots, L\}$, the damage function(s) takes the form

$$\begin{aligned} D(T_{1,t}) &= 1 - \exp[-\xi_0 \exp[\xi_1 T_{1,t}] + \xi_0], \quad \xi_0 \in \mathbb{R} && \text{(basic \& extended ACE),} \\ D_K(T_{1,t}) &= 1 - \exp[-\bar{\xi}_0 \exp[\xi_1 T_{1,t}] + \bar{\xi}_0], \quad \bar{\xi}_0 \in \mathbb{R} && \text{(extended ACE, section 2.2),} \end{aligned}$$

the mean in the equation of motion (6) for temperature layer $i \in \{1, \dots, l\}$ takes the form

$$\begin{aligned} \mathfrak{M}_i^\sigma(T_{i,t}, w_i^{-1} T_{i-1,t}, w_{i+1} T_{i+1,t}) &= \frac{1}{\xi_i} \log \left((1 - \sigma_{i,i-1} - \sigma_{i,i+1}) \exp[\xi_i T_{i,t}] \right. \\ &\quad \left. + \sigma_{i,i-1} \exp[\xi_i w_i^{-1} T_{i-1,t}] + \sigma_{i,i+1} \exp[\xi_i w_{i+1} T_{i+1,t}] \right), \end{aligned} \quad (7)$$

and the parameters ξ_i take the values $\xi_1 = \frac{\log 2}{s} \approx \frac{1}{4}$ and $\xi_{i+1} = w_{i+1} \xi_i$ for $i \in \{1, \dots, l-1\}$ (with w_i , $i \in \{1, \dots, l-1\}$, given).

The coefficients φ in the value function are the shadow values of the respective state variables, and $^\top$ denotes the transpose of a vector of shadow values. The coefficient vector on the resource stock, $\boldsymbol{\varphi}_{R,t}^\top$, has to be time-dependent: the shadow values of the exhaustible resources increases over time following the endogenously derived Hotelling rule. The process φ_t captures the value contribution of the exogenous processes, including technological progress. The damage function is of a double-exponential form with a free parameter ξ_0 , which scales the severity of damages at a given temperature level. The damage parameter ξ_0 is the semi-elasticity of net production with respect to a change of transformed atmospheric temperature $\tau_{1,t} = \exp(\xi_1 T_{1,t})$. The generalized mean \mathfrak{M}_i^σ uses the non-linear weighting function $\exp[\xi_i \cdot]$. Section 2.5 shows that these assumptions match the actual climate dynamics and current assumptions about economic damages. It calibrates the weight matrix $\boldsymbol{\sigma}$, the atmosphere-ocean equilibrium temperature difference w_1 , and the damage parameter ξ_0 .

Expressed in terms of the vector of transformed temperature states $\boldsymbol{\tau}$, the temperatures' equations of motion (7) take the linear form

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1.$$

The parameter σ^{forc} is the weight on radiative forcing in the atmospheric temperature’s equation of motion. It determines the speed of the (initial) response of atmospheric temperature to the greenhouse effect. Under the assumptions of Proposition 1, the optimal consumption rate is $x_t^* = 1 - \beta\kappa$. Society consumes less the higher the discounted shadow value of capital ($x_t^* = \frac{1}{1+\beta\varphi_\kappa}$ with $\varphi_\kappa = \frac{\kappa}{1-\beta\kappa}$), resulting in a consumption rate that decreases in the capital share of output κ . The other controls depend on the precise form of the energy sector.

2.5 Calibration

This section summarizes the model’s base calibration with an emphasis on the novel ocean-atmosphere temperature system and the damage function. I use the carbon cycle of DICE 2013, the a capital share of $\kappa = 0.3$, and the International Monetary Fund’s (IMF) 2015 investment rate forecast of $1 - x^* = 25\%$ to calibrate the annual rate of pure time preference to $\rho = 1.75\%$. Present world output Y is ten times (time step) the IMF’s global economic output forecast of $Y_{2015}^{annual} = 81.5$ trillion USD.

Economic damages are crucial and yet hard to determine. The most wide-spread IAM DICE uses the damage function $D(T) = \frac{1}{1+0.0028T^2}$. Nordhaus (2008) calibrates the coefficient 0.0028 based on a damage survey for a 2.5C warming. I calibrate ACE’s damage coefficient to match Nordhaus’ calibration points of 0 and 2.5°C exactly, delivering the damage semi-elasticity $\xi_0 = 0.0222$. Figure 2 compares the resulting damage curve to that of the DICE-2007 model. The figure also depicts the damage curve $D(T) = 1 - 1 / \left(\left(1 + \frac{T}{20.46}\right)^2 + \left(\frac{T}{6.081}\right)^{6.754} \right)$ suggested by Weitzman (2010), who argues that little is known about damages at higher temperature levels, and that damages might turn out much more convex at high temperatures than assumed in DICE. As compared to DICE-2007, the base calibration of ACE’s damage function generates slightly higher damages below a 2.5°C warming, slightly lower damages above a 2.5°C until a warming of 12°C warming, and higher damages at a warming above 12°C, implying a hard-to-conceive change of life on the planet. Figure 2 also depicts two dashed variations of ACE’s damage function. The lower curve reduces the damage parameter by 50%, resulting in a damage function that lies almost everywhere below DICE. The higher curve increases the damage parameter by 50%, resulting in a damage function that lies everywhere above that of DICE. The analytic solution permits a simple evaluation of such changes in the damage calibration.

The calibration of temperature dynamics (equation 7) uses the Representative Concentration Pathways (RCP) of the recent assessment report by the Intergovernmental Panel on Climate Change IPCC (2013). I use the Magicc6.0 model by Meinshausen et al. (2011) to simulate the RCP scenarios over a time horizon of 500 years. The model emulates the results of the large atmosphere-ocean general circulation models (AOGCMs) and is employed in the IPCC’s assessment report. DICE was calibrated to a (single) scenario using an earlier version of Magicc. My calibration of ACE uses three ocean layers (upper, middle, and deep)

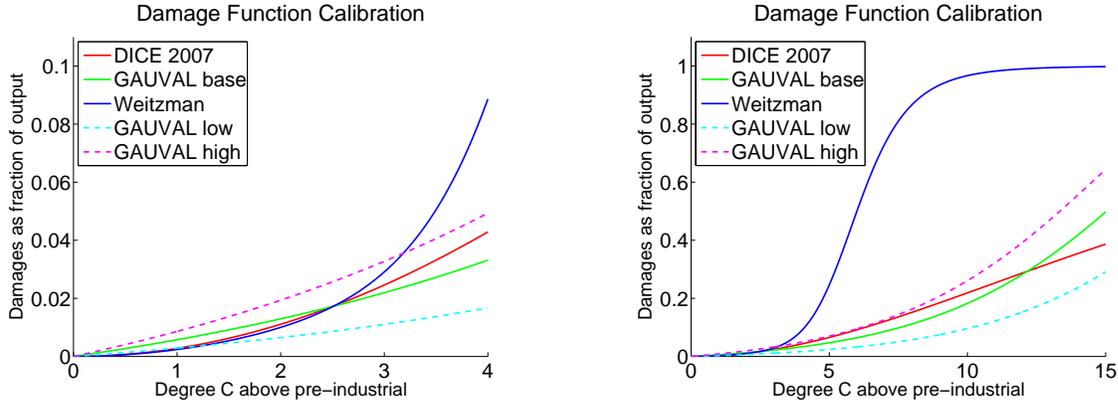


Figure 2: ACE’s damage function compared to that of DICE-2007 and a highly convex damage function suggested by Weitzman (2010). All three lines coincide for a 2.5°C warming, the common calibration point based on Nordhaus (2008). The dashed curves depict ACE’s damage function for a $\pm 50\%$ variation of the base case damage coefficient $\xi_0 \approx 0.022$.

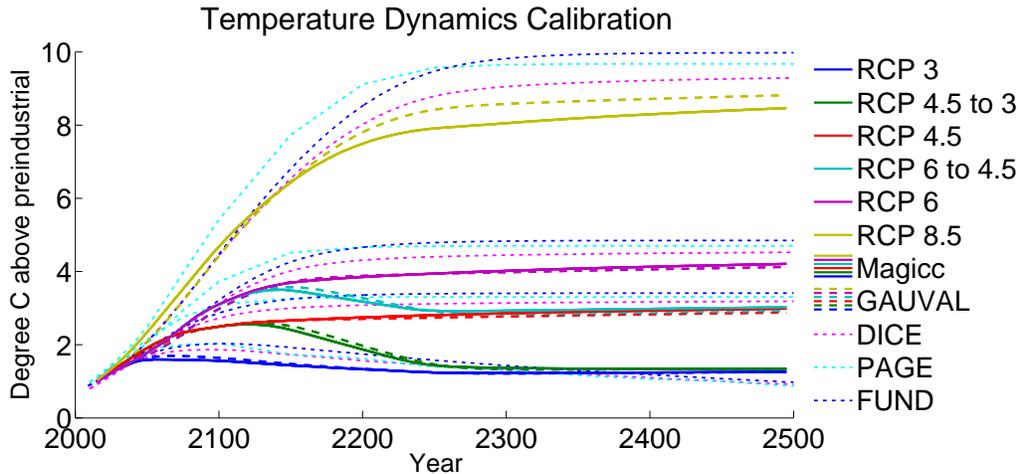


Figure 3: ACE’s response vis a vis Magicc’s response to the color coded radiative forcing scenarios used in the latest IPCC assessment report. RCP 3 is the strongest stabilization scenario and RCP 8.5 is a business as usual scenario. The Magicc model (solid lines) emulates the large atmosphere-ocean general circulation models and is used in the IPCC’s assessment reports. ACE (dashed lines) matches Magicc’s temperature response very well for the “moderate” warming scenarios where the lines hardly differ and reasonably well for RCP 8.5. By courtesy of Cael & Stainforth (2015) the figure presents as well the corresponding temperature response of DICE 2013, PAGE 09, and FUND 3.9, the numeric IAMs used for the interagency report determining the official SCC in the United States. ACE performs better in all scenarios.

compared to Magicc’s 50 and DICE’s single ocean layer(s).

Figure 3 shows the calibration results. The solid lines represent Magicc’s response to the radiative forcing of the RCP scenarios (benchmark), whereas the dashed lines represent ACE’s atmospheric temperature response. In addition to the original RCP scenarios, I include two scenarios available in Magicc6.0 that initially follow a higher radiative forcing scenario and then switch over to a lower scenario (RCP 4.5 to 3 and RCP6 to 4.5). These

scenarios would be particularly hard to fit in a model tracing only atmospheric temperature. The ability to fit temperature dynamics across a peak is important for optimal policy analysis. ACE’s temperature model does an excellent job in reproducing Magicc’s temperature response for the scenarios up to a radiative forcing of $6\text{W}/\text{m}^2$. It performs slightly worse for the high business as usual scenario RCP8.5, but still well compared to other IAMs.

3 Results from the Deterministic Model

The social cost of carbon (SCC) is the money-measured present value welfare loss from an additional ton of CO_2 in the atmosphere. The economy in section 2.1 decentralizes in the usual way and the Pigovian carbon tax is the SCC along the optimal trajectory of the economy. In the present model, the SCC is independent of the future path of the economy and, thus, this unique SCC is the optimal carbon tax. The present section discusses the interpretation and quantification of its closed-form solution. It explores the social cost of global warming and the social benefits of carbon sequestration. A proposition establishes that mass conservation in the carbon cycle makes the SCC highly sensitive to pure time preference (and not to the consumption discount rate in general).

3.1 The Price of Atmospheric Carbon

Appendix C.3 solves for the shadow values and derives the optimal CO_2 tax. It is proportional to output Y_t and increases over time at the rate of economic growth as in Golosov et al. (2014). ACE avoids the infinite sum of an impulse response formulation and derives in detail how climate dynamics impacts the optimal tax.

Proposition 2 (1) *Under the assumptions of section 2, the SCC of the basic ACE model in money-measured consumption equivalents (USD 2015) is*

$$SCC_t = \underbrace{\frac{\beta Y_t}{M_{pre}}}_{25.5 \frac{\$}{tC}} \underbrace{\xi_0}_{2.2\%} \underbrace{[(1 - \beta \sigma)^{-1}]_{1,1}}_{1.4} \underbrace{\sigma^{forc}}_{0.42} \underbrace{[(1 - \beta \Phi)^{-1}]_{1,1}}_{3.7} = 56.5 \text{ \$/tC} \quad (8)$$

where $[\cdot]_{1,1}$ denotes the first element of the inverted matrix in squared brackets, and the numbers rely on the calibration discussed in section 2.5.

The ratio of production to pre-industrial carbon concentrations M_{pre} sets the units of the carbon tax. The discount factor β reflects a one period delay between temperature increase and production impact. The damage parameter ξ_0 represents the constant semi-elasticity of net production to a transformed temperature increase, i.e., to an increase of $\tau_1 = \exp(\xi_1 T_1)$. In the absence of any interesting climate dynamics, these terms would imply a carbon tax of 25.5\$ per ton of carbon.

A von Neumann series expansion of the (bounded operator) $\beta\Phi$ helps to interpret the terms in square bracket governing carbon and temperature dynamics, e.g.,

$$(1 - \beta\Phi)^{-1} = \sum_{i=0}^{\infty} \beta^i \Phi^i .$$

The element $[\Phi^i]_{1,1}$ of the transition matrix characterizes how much of the carbon injected into the atmosphere in the present remains in or returns to the atmospheric layer in period i , after cycling through the different carbon reservoirs. E.g., $[\Phi^2]_{1,1} = \sum_j \Phi_{1,j} \Phi_{j,1}$ states the fraction of carbon leaving the atmosphere for layers $j \in \{1, \dots, m\}$ in the first time step and arriving back to the atmosphere in the second time step. Thus, the term $[(1 - \beta\Phi)^{-1}]_{1,1}$ characterizes in closed form the discounted sum of CO₂ persisting in and returning to the atmosphere in all future periods. The discount factor accounts for the delay between the act of emitting CO₂ and the resulting temperature forcing over the course of time. Quantitatively, the persistence of carbon increases the earlier value of 25.5\$/tC by a factor of 3.7. The resulting carbon tax would be 95\$/tC – if ignoring temperature dynamics.

The terms $[(1 - \beta\sigma)^{-1}]_{1,1} \sigma^{forc}$ capture the atmosphere-ocean temperature dynamics resulting in both delay and persistence. Analogously to the interpretation in the case of carbon, the expression $[(1 - \beta\sigma)^{-1}]_{1,1}$ characterizes the generalized heat flow that enters, stays, and returns to the atmospheric layer. Thus, the simple closed-form expression for the carbon tax in equation (8) captures an infinite double-sum: an additional ton of carbon emissions today causes radiative forcing in all future periods, and the resulting radiative forcing in any given period causes warming in all subsequent periods. The parameter σ^{forc} captures the speed at which the atmospheric temperature responds to radiative forcing. The response delay, a factor around 0.4, substantially reduces the SCC. However, at the same time, the ocean implied temperature persistence increases the SCC by a factor of 1.4. Together, the ocean-atmosphere temperature dynamics reduce the carbon tax by a factor of 0.6 resulting in the value of 56.5 USD per ton of carbon.⁷

Expressed in tons of CO₂, this SCC is 15.5 USD, coinciding up to one dollar to the DICE-2007 carbon tax for 2015.⁸ At the gas pump, the SCC translates into 14 cent per gallon or 4 cent per liter. The (dashed) variation of the damage function in Figure 2 implies a $\pm 50\%$ variation of the semi-elasticity ξ_0 and, thus, the SCC. Ignoring the transitory atmosphere-ocean temperature dynamics calibrated in Figure 3 would overestimate the carbon tax by

⁷Golosov et al. (2014) and Gerlagh & Liski (2012) use an emission response model mostly equivalent to the common carbon cycle models that I adopt here. Their models do not explicitly incorporate radiative forcing, temperature dynamics, and damages as a function of temperature. However, Gerlagh & Liski (2012) introduce a reduced form damage delay component that gets at the important delay between peak emissions and peak damages. This delay multiplier contributes a factor of .45 in their closest scenario (“Nordhaus”), which cuts the tax a little more than ACE’s factor of $1.4 \cdot 0.42 \approx .6$ based on an explicit model of temperature dynamics.

⁸DICE expresses the carbon tax in 2015 in USD of 2005, which have to be translated into 2015 USD for the comparison.

70%. Ignoring carbon persistence would result in a carbon tax that is only 27% of its actual value.

The SCC in equation (8) is independent of the atmospheric carbon concentration and of the prevailing temperature level. A corresponding independence already prevails in Golosov et al. (2014), and it opposes the common perception that slacking on climate policy today will require more mitigation in the future. This result might sound like good news, but what the model really says is: if we delay policy today, we will not compensate in our mitigation effort tomorrow, but live with the consequences forever. Yet, the result contains some good news for policy makers and modelers. Setting the optimal carbon tax requires minimal assumptions about future emission trajectories and mitigation technologies. The policy maker sets an optimal price of carbon and the economy determines the resulting optimal emission trajectory. The result also explains the generality of equation (2), and why some seemingly important parameters in numeric models hardly influence the optimal tax. These findings connect immediately to the debate on the slope of the marginal damage curve in the “taxes versus quantities” literature (Weitzman 1974, Hoel & Karp 2002, Newell & Pizer 2003). ACE states that the marginal social damage curve for CO₂ emissions is flat. In consequence, taxes eliminate the welfare cost under technological uncertainty and asymmetric information, and the policy maker does not have to adjust the tax to technology shocks from period to period.

The common intuition that the SCC ought to increase in the CO₂ concentration and the prevailing temperature level results from the convexity of damages in temperature. ACE clearly has such a convex damage function (Figure 2) and yet marginal damages from releasing a ton of CO₂ are independent of the emission level, the CO₂ concentration, and the temperature level. The main (but not the only) reason lies in the radiative forcing equation (5): the higher the CO₂ concentration, the less does an additional ton of emissions contribute to further forcing and, thus, warming. A simple intuition is that CO₂ traps (absorbs) a certain spectrum of the wavelength that our planet radiates out into space, thereby warming the planet. If there is already a high concentration of CO₂ in the atmosphere, most of the energy leaving the planet in this wavelength is already trapped. An additional unit of CO₂ emissions has a lower warming impact than the earlier units. In addition, the prevailing temperature level does not affect the CO₂ because a higher temperature level at a given CO₂ concentration implies less warming in the future. Whether the “saturation” effect in CO₂ absorption and the temperature dynamics indeed align with the convenient formula of ACE’s SCC formula is an empirical question. ACE translates these assumptions into (i) the shape of a damage function on temperature that cancels the logs of radiative forcing and (deterministic) welfare, and (ii) a novel functional form and calibration of the ocean-atmosphere temperature dynamics. Figures 2 and 3 state that, indeed, damages seem reasonable enough given the current state of knowledge, and temperature dynamics are surprisingly well aligned with the model’s assumptions.

3.2 The Role of Calibration, Discounting, and Interactions

The formula for the SCC derived in the preceding section directly lays out its dependence on world output and damages. It also explained the basic dependence on carbon and temperature dynamics and its interaction with time preference. This section takes a more careful look at this interaction and the role of discounting. It also compares the SCC between the basic ACE (Ramsey) and the extended (endogenous growth) model. Finally, I present an alternative formula for the social cost of carbon that does not depend on the endogenous world output, which itself responds to climate policy. The formal results are summarized in the following proposition.

Proposition 3 (1) *A carbon cycle (equation 4) satisfying mass conservation of carbon implies a factor $(1 - G\beta)^{-1}$, approximately proportional to $\frac{1}{\rho}$, in the closed-form solution of the SCC (equation 8).*

(2) *In money-measured capital equivalents (USD 2015) the SCC is*

$$SCC_t^{*,basicACE} = \frac{K_t}{\kappa} \xi_0 \left[(1 - \beta\sigma)^{-1} \right]_{1,1} \frac{\beta\sigma^{forc}}{M_{pre}} \left[(\mathbf{1} - \beta\Phi)^{-1} \right]_{1,1}$$

$$SCC_t^{*,extendedACE} = \frac{K_t}{\kappa} \left((1 - \beta) \xi_0 + \beta\kappa\bar{\xi}_0 \right) \left[(1 - \beta\sigma)^{-1} \right]_{1,1} \frac{\beta\sigma^{forc}}{M_{pre}} \left[(\mathbf{1} - \beta\Phi)^{-1} \right]_{1,1}.$$

(3) *Let all damage coefficients coincide within and between the two models: $\xi_0^{basic} = \xi_0^{extended} = \bar{\xi}$. Then, for the range of reasonable model calibrations,⁹ it is $SCC_t^{basicACE} < SCC_t^{extendedACE}$ when expressing the social cost in consumption units and $SCC_t^{*,basicACE} > SCC_t^{*,extendedACE}$ when expressing the social cost in capital units.*

It is well-known that the consumption *discount rate* plays a crucial role in valuing long-run impacts. Finding (1) is different. It states that the interaction of *pure time preference* and carbon cycle dynamics is the main sensitivity when it comes to discounting. The consumption discount rate is the sum of pure time preference and growth based discounting (higher future wealth reducing marginal utility from a unit of future consumption). Finding (1) sheds new light on a debate between Stern (2007) and (Nordhaus 2007). The Stern Review of Climate Change argues in favor of a vanishingly small rate of pure time preference on normative grounds and finds an unusually high SCC. (Nordhaus 2007) argues that such a choice of the pure rate of time would not make a difference if only the Stern Review had adjusted the growth part of the discount rate to match observed rates. Yet, Proposition 3's part (1)

⁹Let κ^b denote the capital share in the basic ACE model and let κ^e denote the capital share in the final goods sector of the extended ACE model. Then, the cited inequalities hold if and only if $\kappa^e < \frac{(1-\beta)\kappa^b}{1-\beta\kappa^b}$. In the present paper's calibration with $\kappa^b = .3$ and $\rho \leq 1.75\%$ a sufficient condition is that the capital intensity in the final goods sector is $\kappa^e > 10\%$. Stacking the cards against the condition, an overall capital share $\kappa^b = .4$ and a rate of pure time preference $\rho = 3\%$ would still require a capital share in the final goods sector κ^e below 15% to flip the sign of the cited inequalities.

points out that the main sensitivity is indeed pure time preference and not growth-based discounting.

The composition of the social discount rate has crucial implications for optimal climate change mitigation. Thus, it is important to *reconsider the standard calibration procedure in integrated assessment*, which assumes that a representative agent with an infinite planning horizon invests in a market giving her the average observed return with certainty. First, Schneider et al. (2013) show in a continuous time overlapping generations model how the common infinitely-lived-agent based calibration of IAMs overestimates the rate of pure time preference under limited altruism. Second, decades of asset pricing observation and theory tell us that the risk-free interest (or discount) rate is far from the average market return. Bansal et al. (2012) calibrate the rate of pure time preference to $\rho = 0.11\%$, carefully disentangling risk attitude and risk premia from consumption smoothing and the risk-free discount rate. Their model explains observed asset prices substantially better than any asset pricing approach based on the standard economic model with higher time preference. Traeger (2012a) shows how uncertainty-based discounting of an agent whose risk aversion does not coincide with her consumption smoothing preference (falsely) manifests as pure time preference in the economic standard model, and discusses implications for climate change evaluation.

The descriptive approach by Bansal et al. (2012), the Stern Review’s normative reasoning tied to a dual role of individuals taking consumption and large-picture policy decisions (Hepburn 2006), and Schneider et al.’s (2013) approach to life cycle-based time preference calibration all lead to a pure rate of time preference around 0.1%, which delivers

$$SCC_t = \underbrace{\frac{\beta Y_t}{M_{pre}}}_{\substack{2.2\% \\ 26 \frac{\$}{tC} \quad 30 \frac{\$}{tC}}} \underbrace{\xi_0}_{1.42} \underbrace{[(1 - \beta \sigma)^{-1}]_{1,1}}_{0.42} \underbrace{\sigma^{forc} [(1 - \beta \Phi)^{-1}]_{1,1}}_{3.726} = 57660 \frac{\$}{tC} .$$

The formula emphasizes that, under certainty, the main determinant of the optimal carbon tax is the interaction of time preference and carbon dynamics, delivering a factor 7 increase by itself. The intuition connects closely to a discussion in Section 3.1. First, the optimal policy asks us to live with the consequences of historic overindulgence in carbon instead of making up by additional mitigation efforts in the future. Second, carbon does not decay. Carbon only cycles through the different reservoirs.¹⁰ It is pure time preference that matters rather than consumption discounting because damages grow proportional to output. The

¹⁰The fact that some of it eventually turns into limestone is negligible for even very long human planning horizons. A comparison of scientific carbon cycle models finds that on average 18% of a 100Gt carbon emission pulse, approximately 10 years of present CO₂ emissions, still remain in the atmosphere after 3000 years (Joos et al. 2013). In DICE 2013’s carbon cycle adopted here, 6% of an anthropogenic emission unit stays in the atmosphere forever, which is calculated as follows. The maximal eigenvalue of the transition matrix Φ is unity. The corresponding eigenvector governs the long-run distribution as the transitions corresponding to all other eigenvectors are damped. I obtain the 0.06 as the first entry of the corresponding eigenvector.

important take-away is twofold. First, the composition of the consumption discount rate matters. Second, discounting related model calibration matter hugely because of carbon persistence.

ACE’s functional forms flesh out and explain a particularly important discounting sensitivity. Several changes in functional forms can reduce the sensitivity of $\frac{1}{\rho}$. Two recent studies derive approximate formulas for the SCC that show a decreasing sensitivity as we move away from logarithmic utility (van den Bijgaart et al. 2016, Rezai & der Ploeg 2016). These changes simultaneously imply that we clean up more of our historic sins, explaining why we worry less about the long-run. More generally, the sensitivity on pure time preference depends on the combination of the intertemporal elasticity of substitution and the elasticity between damages and output. A different change can make the model even more sensitive to pure time preference. Assume a growing population at rate g instead of normalizing population to unity. Then, the common welfare measure using population-weighted per capita consumption results in the factor $\frac{1}{\rho-g}$ rather than $\frac{1}{\rho}$, making the optimal carbon tax even more sensitive to low rates of pure time preference. The intuition for this result is of more general interest to understand how population growth affects the optimal carbon tax in integrated assessment models: population weighted per-capita consumption puts additional weight on future generations that are more numerous, acting as a reduction of time preference. Finally, I note that the temperature dynamic system does not give rise to a similar sensitivity because heat is constantly exchanged with outer space (no hysteresis). Appendix C.1 illustrates Proposition 3’s Part (1) for a two-layer carbon cycle, and Appendix C.2 the absence of such sensitivity for a two-layer atmosphere-ocean temperature system. It also shows how a frequently used decay approximation of the carbon cycle misses the sensitivity to pure time preference.

Part (2) of Proposition 3 spells out a different representation of the optimal carbon tax. Imposing an optimal carbon tax affects world output. Therefore, earlier closed-form approaches such as Golosov et al. (2014) and Gerlagh & Liski (2012) have been criticized for their reliance on world output in the SCC’s formula. The concern is that the tax value calculated before imposing the tax will no longer be optimal after imposing the tax because world output changes in response to the tax. Most IAMs suggest an output response in the low percentage order when switching from business as usual to optimal policy, making it more of a theoretical concern of internal model consistency than of quantitative importance. Part (2) of Proposition 3 shows how to prevent the concern by presenting the SCC in units of the pre-determined capital stock. Quantitatively, the calibration gives similar results as $\frac{K_t}{\kappa} \approx Y_t$. More precisely, the estimates of κ and the capital stock are usually entangled, and differences would emerge mostly because models used to estimate κ and capital can deviate from the assumed Cobb-Douglas specification.

Expressing the SCC in capital units is most interesting to compare the SCC deriving from the basic ACE model to that of the extended ACE. Part (2) of Proposition 3 shows that

the endogenous growth specification with the dedicated capital sector replaces the damage parameter ξ_0 by the factor $(1 - \beta)\xi_0 + \beta\kappa\bar{\xi}_0$. In addition, κ in the extended ACE merely characterizes the capital intensity of the final consumption goods sector, which is lower than the overall capital intensity of the economy represented by κ in the basic ACE model. In both models, ξ_0 specifies the fractional loss of final output in response to a temperature increase. In the basic ACE model final output affects both consumption and investment. In the extended ACE model of section 2.2, ξ_0 only governs consumption and the separate parameter $\bar{\xi}_0$ governs the damages to the capital stock. To understand the formula, note that $1 - \beta$ is the weight on the present period, in which consumption damages arise, and β is the weight on the future affected through capital damages. In addition, the damages to the capital stock translate into welfare relevant consumption loss only through the consumption sector's capital value share κ .

Part (3) of Proposition 3 further examines the case where damages in the consumption and the capital sector coincide. Then, the SCC of the extended model is larger than the SCC resulting from the basic ACE when using consumption equivalents to obtain the USD value. This relation flips if we use capital units to convert the welfare loss into USD equivalents. The reason is that the endogenous growth AK version gives rise to a higher shadow value of capital. A higher shadow value of capital implies a larger welfare loss from a given temperature increase. Using present consumption for value conversion implies a higher optimal carbon tax. Yet, if we use capital units as the basis for value conversion, then the consumption damages weigh less in the extended ACE as they are converted into the more valuable capital units. Karp & Rezai (2017) point out the importance of non-substitutability between consumption and capital units to understand the relative value changes that distinguish winners and losers (in particular in their OLG economy). The extended ACE permits such a distinction and points out that different ways to calibrate the Ramsey versus AK model results in different relative rankings of the SCC.

4 Uncertainty: the Stochastic ACE

The occurrence of climatic change is no longer uncertain. Data over the past centuries and over pre-historic time scales suggest a strong correlation between climate and atmospheric carbon dioxide. Basic causal interactions including carbon dioxide's absorption of outgoing radiation (greenhouse effect) can be measured in the laboratory. Yet, the warming that results from a given emission trajectory is highly uncertain. First, we have a limited understanding of how carbon dioxide builds up in the atmosphere and, second, the temperature response to an increase in atmospheric carbon dioxide is highly uncertain. This section extends ACE to incorporate uncertainty. In particular, it characterizes a general class of stochastic processes for which ACE permits analytic solutions.

4.1 The Main Climate Uncertainties

Over 10% of the annual flow of *anthropogenic carbon* emissions leave the atmosphere into an unidentified sink. These missing 1Gt+ in the carbon budget are over twice the weight of all humans walking the planet. Current research is not conclusive, but a likely candidate for at least part of the “missing sink” are boreal or tropical forests. Most importantly, our lack of understanding current carbon flows implies major uncertainties in predicting future carbon dynamics and whether e.g. the missing sink uptake is permanent or temporary. Society invests substantial sums into the reduction of these uncertainties, including the launching of satellites and new supercomputing facilities. ACE can produce a simple estimate of the welfare costs of these uncertainties, serve as a formal model for quantifying the benefits of uncertainty reduction, and analyze the implications for the optimal mitigation policy.

Even if we knew future atmospheric carbon concentrations, we remain uncertain about the implied warming. The deterministic ACE assumed that a doubling of the pre-industrial CO₂ concentration from 280 ppm to 560 ppm yields a medium-term temperature increase of 3C (climate sensitivity). At present, CO₂ levels are up to almost 400 ppm. Including the CO₂ equivalent forcing of other GHGs, the level is already close to 480 ppm. The present warming is still much below the corresponding equilibrium increase because of the atmosphere-ocean temperature interaction discussed in sections 2.3 and 3.1. The assumed 3C warming is a guesstimate. The value depends on a set of uncertain feedback processes that either increase or decrease the initial warming. A simple feedback example is that higher temperatures imply more evaporation, and water vapor itself is a powerful GHG. The value of 3C was cited as the best guess in the first four IPCC assessment reports. The latest report deleted this best guess and only cites a likely range of 1.5-4.5C (IPCC 2013), where likely characterizes a 66% probability interval.

I suggest two conceptually different ways to think about these uncertainties. In the first interpretation, the decision maker worries merely about (unpredicted) changes in the carbon flows or temperature feedbacks over time. A small persistent shock moves the carbon or heat flow and either increases or decreases sink uptake or temperature. Over time, these shocks accumulate and so does the uncertainty in forecasting future carbon levels, temperatures, and economic damages. In the second interpretation, the main uncertainty is epistemological: it reflects a lack of knowledge in the scientific community. The present decision maker is uncertain about the dynamics, possibly including the variance and skewness of the underlying distributions. This uncertainty prevails today, but we expect to acquire more knowledge about the true underlying dynamics in the future, making the system increasingly predictable. I start out with a generic treatment of uncertainty until explicitly distinguishing between the two conceptual differences in section 5.2.

4.2 Risk Attitude

Logarithmic utility provides a reasonable description of intertemporal substitutability. However, the assumption performs poorly in capturing risk attitude. The long-run risk literature estimates the coefficient of relative risk aversion of a representative household closer to 10 than to unity (Vissing-Jørgensen & Attanasio 2003, Bansal & Yaron 2004, Bansal et al. 2010, Chen et al. 2013, Bansal et al. 2012).¹¹ Merely increasing the utility function’s curvature would result in high risk-free discount rates that cannot be reconciled with market observation (risk-free rate puzzle) and would unwarrantedly discount away worries about the future climate. Moreover, the market rejects the assumption that the intertemporal elasticity of substitution fully determines risk attitude, which is an assumption built into the intertemporally additive expected utility (standard) model but is not implied by the von Neumann & Morgenstern (1944) axioms. I follow the asset pricing literature, an increasing strand of macroeconomic literature, and some recent numeric approaches to climate change assessment in using Epstein-Zin-Weil preferences. This approach accommodates a realistic coefficient of risk aversion, disentangling it from the unit elasticity of intertemporal substitution.

I denote the underlying probability space by $(\Omega, \mathcal{F}, \mathbb{P})$. Some scenarios require explicit informational states denoted by the vector \mathbf{I}_t . The Bellman equation under uncertainty is

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, \mathbf{I}_t, t) = \max_{x_t, \mathbf{N}_t} \log c_t \tag{9}$$

$$+ \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha (V(k_{t+1}, \boldsymbol{\tau}_{t+1}, \mathbf{M}_{t+1}, \mathbf{R}_{t+1}, \mathbf{I}_{t+1}, t)) \right] \right).$$

Expectations \mathbb{E}_t are conditional on time t information.¹² The non-linear uncertainty aggregator is a generalized mean $f^{-1} \mathbb{E}_t f$ with $f(\cdot) = \exp(\alpha \cdot)$. A positive parameter α characterizes intrinsic risk loving, and a negative parameter characterizes intrinsic risk aversion.

Epstein & Zin’s (1991) original definition of disentangled Arrow-Pratt risk aversion delivers the coefficient of constant relative risk aversion $\text{RRA} = 1 - \frac{\alpha}{(1-\beta)}$. In the present model, it is not Arrow-Pratt risk aversion that drives risk averse behavior, but the parameter $-\alpha$ itself. Intuitively, $-\alpha$ measures how much more averse a decision maker is to risk than to deterministic consumption fluctuations. The limit $\alpha \rightarrow 0$ recovers the usual Bellman equation where risk aversion is merely generated by aversion to intertemporal inequality. Traeger (2014) gives an axiomatic definition of α as an intrinsic measure of risk attitude. The asset pricing literature estimates Epstein & Zin’s (1991) Arrow-Pratt risk aversion parameter RRA in the range of [6, 9.5], which implies $\alpha \in [-1.5, -1]$. Figure 6 in Appendix F.1 illustrates

¹¹Nakamura et al. (2013) obtain one of the lowest estimates by combining the long-run risk model and the Barro-Riesz model, still resulting in a coefficient of relative risk aversion of 6.4.

¹²The space will be equipped with the filtration $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$ generated by the stochastic processes driving carbon accumulation and temperature. The filtration \mathcal{F}_t captures all information available at time t and $\mathbb{E}_t(\cdot) \equiv \mathbb{E}(\cdot | \mathcal{F}_t)$. The state vector \mathbf{I}_t captures the structural information affecting the value function.

the corresponding risk aversion for a small and a large binary lottery. The analytic formulas make it easy for the reader to vary the degree of risk aversion for the quantitative results.

4.3 Stochastic Motion and General Solution

I focus on the uncertainty induced by the stochastic evolution of carbon \mathbf{M}_t and temperature $\boldsymbol{\tau}_t$. The scenarios incorporating structural uncertainty or persistent shocks require an additional equation of motion for their informational variables

$$\mathbf{I}_{t+1} = h \left(\Phi, \boldsymbol{\sigma}, \mathbf{M}_t, \boldsymbol{\tau}_t, \mathbf{I}_t, \sum_{i=1}^{I^d} E_{i,t} + E_t^{exo}, \frac{\sigma^{forc}}{M_{pre}}, G_t, \omega \right),$$

where updating the information depends on the realization $\omega \in \Omega$.

The proposition below characterizes a class of stochastic processes that permit an analytic solution of the stochastic ACE model. A solution can either result in an explicit formula for the optimal carbon tax or it can define the optimal policy implicitly by translating the dynamic optimization problem into a set of simultaneous algebraic equations. The proposition relies on the following observations. First, an affine value function solves the deterministic model. Second, if the value function is affine, then the expectation formation in the Bellman equation (9) resembles a cumulant generating function $G_{\mathbf{X}}(\mathbf{z}) \equiv \log [\mathbb{E} \exp(\mathbf{z}\mathbf{X})]$ of a random vector \mathbf{X} . The cumulant generating function is the logarithm of the moment generating function. Third, if we find a cumulant generating function that preserves the value function's affine structure, then the same procedure that solves the deterministic model promises to solve the stochastic model.

Proposition 4 *Let $\mathbf{X}_t = (\mathbf{M}_t, \boldsymbol{\tau}_t, \mathbf{I}_t) \in \mathbb{R}^N$ follow an affine process whose conditional cumulant generating function satisfies*

$$G_{\mathbf{X}_{t+1}}(\mathbf{z}) = \log [\mathbb{E} (\exp(\mathbf{z}\mathbf{X}_{t+1}) | \mathbf{X}_t)] = a(\mathbf{z}) + \sum_{i=1}^N b_i(\mathbf{z}) X_{t,i}. \quad (10)$$

Then, an affine value function solves ACE's dynamic programming problem if and only if the set of shadow values $\boldsymbol{\varphi}_M^\top, \boldsymbol{\varphi}_\tau^\top, \boldsymbol{\varphi}_I^\top$ solves the algebraic equations

$$\begin{aligned} \varphi_{M,i} &= \frac{\beta}{\alpha} b_i^M(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top) & \forall i = 1, \dots, m \\ \varphi_{\tau,i} &= \frac{\beta}{\alpha} b_i^\tau(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top) - \delta_{i,1}(1 + \beta \varphi_k) \xi_0 & \forall i = 1, \dots, l \\ \varphi_{I,i} &= \frac{\beta}{\alpha} b_i^I(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top) & \forall i = 1, \dots, N - l - m, \end{aligned}$$

where $(b_1^M, \dots, b_m^M, b_1^\tau, \dots, b_l^\tau, b_1^I, \dots, b_{N-l-m}^I) = (b_1, \dots, b_N)$ and $\delta_{i,j}$ denotes the Kronecker-delta (one if $i = j$ and zero otherwise). The shadow value $\varphi_{M,1}$ determines the optimal carbon tax. The function $a(\mathbf{z})$ does not affect the optimal policies. It directly affects welfare.

Proposition (11) transforms a high dimensional and difficult-to-solve stochastic dynamic optimization problem on the state space into a simple root finding problem for the shadow values (co-states), which solves either in closed form or trivially on a computer.

Most analytic assessments of climate change under uncertainty have focused on the welfare impact under uncertainty. Proposition (11) points out that even major welfare losses may not change the optimal carbon tax if the uncertainty operates through the affine part $a(\mathbf{z})$ of equation (11). The impact of uncertainty on the optimal policy depends on its interaction with the state variables. The set of stochastic processes satisfying equation (11) is large and includes the autoregressive shock model with almost arbitrary distributions, the normal-normal Bayesian learning model, the Gaussian square root process, and the autoregressive gamma model (Gourieroux & Jasiak 2006, Le et al. 2010).

5 Results under Uncertainty

The section discusses policy and welfare change under uncertainty. First, I compare the optimal carbon tax response to uncertainty governing carbon flows versus temperature response, and governing conditional expectations versus stochastic volatility. Second, I contrast stochasticity, epistemological uncertainty, and learning in the welfare context and focus on the contributions of risk aversion, time preference, distributional assumptions, and their interactions.

5.1 Uncertainty and the Optimal Carbon Tax

The present section analyzes the optimal carbon tax under uncertainty. I introduce a stochastic dynamic system that permits a closed-form solution. The model borrows heavily from the long run risk literature in asset pricing. The crucial distinction to that literature is that I endogenize the (here: climate) risk so that there is an interaction between choice (policy) and aggregate risk. The one-step-ahead equation for the carbon stock (formerly equation 4) now becomes

$$\mathbf{M}_{t+1} = \Phi \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1 + \mathbf{x}_t^M + \boldsymbol{\sigma}_t^M \mu_t^M \quad (11)$$

where \mathbf{x}_t^M denotes persistent deviations of the conditional (one-step-ahead) expectations. The vector $\mathbf{x}_t^M = (1, -1, 0)^\top x_t^M$ captures that atmospheric carbon does not decay but moves between atmosphere (first reservoir) and the combined upper ocean and biosphere sink (second reservoir). The independent white noise shock $\mu_t^M \sim N(0, 1)$ makes $\boldsymbol{\sigma}_t^M$ the conditional (one-step-ahead) variance of carbon flows that depends on the scalar process σ_t^M as $\boldsymbol{\sigma}_t^M = (1, -1, 0)^\top \sigma_t^M$. The equations of motion governing conditional expectations and

variance are

$$x_{t+1}^M = \gamma_M^x x_t^M + \delta_M^{Mx} \sqrt{\frac{M_{1,t}}{M_{pre}} - \eta_M} \chi_t^M + \delta_M^{\sigma x} \sigma_t^M \omega_t^M \quad (12)$$

$$\sigma_{t+1}^{M^2} = \gamma_M^\sigma \sigma_t^{M^2} + \delta_M^{M\sigma} \left(\frac{M_{1,t}}{M_{pre}} - \eta_M \right) + \bar{\sigma}^M \nu_t^M \quad (13)$$

The γ -parameters characterize the persistence of the shocks to the mean (first equation) and the volatility process (second equation). As I show in section 5.2, epistemological uncertainty corresponds to a high persistence. The δ -parameters in the second terms on the right hand side capture the endogeneity of climate risk. The uncertainty of the shocks to the mean and the stochasticity of the conditional variance increase as we deviate further from the pre-industrial equilibrium ($M_{1,t} = M_{pre}$). The η -parameter introduces some freedom in calibrating how quickly uncertainty increases away from the pre-industrial equilibrium. For the stochastic volatility equations, the last term on the right hand side specifies the exogenous uncertainty (prevailing already at pre-industrial times), where $\nu_t^M \sim N(0, 1)$ and $\bar{\sigma}^M$ characterizes the corresponding variance. Finally, the parameter $\delta_M^{\sigma x}$ permits coupling the stochastic volatility to the conditional expectations process, where $\omega_t^M \sim N(0, 1)$ is once more white noise. In summary, equation (13) captures that future expected carbon flows are uncertain, and equation (14) captures that the variance of carbon flows is itself uncertain. In the asset pricing literature, the second type of uncertainty has a crucial impact on asset prices. It remains to be seen how important the role of such uncertainty volatility is in the climate context.

Appendix ?? states the analogous Gaussian uncertainty model for the temperature's equation of motion, as well as the resulting closed-form solution of the carbon tax. This result permits a direct comparison of similarities and differences in the carbon tax adjustment that arise merely from their roles in the impact chain. Yet, such a model cannot capture a crucial difference between carbon flow dynamics and temperature dynamics. Temperature $T_{i,t} = \frac{1}{\xi_i} \log(\tau_{i,t})$ is a logarithmic transformation of the generalized temperature state. The uncertainty model for carbon flows in equation (13) implies that the expected carbon dynamics coincides with the deterministic carbon dynamics (assuming the initial conditional expectation $x_0^M = 0$). Introducing uncertainty into the temperature equation has to hold expected temperature dynamics approximately equal to its deterministic baseline. Otherwise, we would discuss the policy impact of changes in expected temperature dynamics rather than the impact of uncertainty. Moreover, negative realization of generalized temperature $\tau_{i,t}$ would imply nonsensical realization of real temperature.¹³ As a result, generalized tem-

¹³Note that the long-run risk model in general also gives rise to nonsensical negative realizations of the variance $\sigma_{t+1}^{M^2}$. This fact is well-known, yet it is widely used as an approximate model with closed-form solution, assuming the actual calibration of the model makes these realization of second order importance. The issue with temperature is more serious. To keep temperature expectations (log expectations of $\tau_{i,t}$) close to the deterministic evolution the model has to be de-biased. Yet, any realization of $\tau_{i,t} = 0$ would cause an infinitely negative expectation. Hence, the model cannot be de-biased in a meaningful way.

perature has to be governed by a positively skewed distribution with a suitable lower bound. For this purpose, I employ the autoregressive gamma model by Gouriou & Jasiak (2006), also employed to the long-run risk literature in asset pricing by Le et al. (2010) and ?.¹⁴ The one-step-ahead state in the autoregressive gamma process is governed by a gamma distribution whose shape parameter is modulated by a realization of a Poisson distribution. I refer to Appendix 5 for details.

Generalized temperature follows the equation of motion

$$\tau_{t+1} = \sigma \tau_t + \left(\sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} + h(x_{t+1}^\tau - z_{t+1}) \right) \mathbf{e}_1 \quad (14)$$

where x_t^τ follows an autoregressive gamma process with one-step-ahead expectations and variance

$$\mathbb{E} x_{t+1}^\tau = \gamma_\tau^x x_t^\tau + \delta_\tau \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \quad (15)$$

$$\text{Var} x_{t+1}^\tau = c \left[2\gamma_\tau^x x_t^\tau + \delta_\tau \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \right] \quad (16)$$

and z_t is a deterministic process neutralizing the temperature expectations (in degree C) to those of the deterministic model¹⁵

$$z_{t+1} = \gamma_\tau^x z_t + (\delta_\tau - \epsilon(c)) \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \quad (17)$$

The model assumes that in the present $z_0 = x_0^\tau$. Once again, the γ -parameter characterizes persistence of the stochastic deviation, and the δ -parameter and the η -parameter capture the responsiveness to the deviation from the pre-industrial equilibrium. Similarly to δ , also the parameter h scales the importance of the stochastic contributions to temperature dynamics. In the proof of Proposition 5 in Appendix G.1 shows that I can either set $h = 1$ or $\delta = 1$ without loss of generality. Henceforth, I will set $\delta = 1$ and only use the parameter h .

Illustrating the model, I first assume $\epsilon(c) = 0$ (no bias-correction). The temperature contribution proportional to σ^{forc} in equation (G.20) continues to capture the expected baseline impact of radiative forcing in response to a given atmospheric CO₂ concentration $M_{1,t}$ (and the forcing of other GHGs measured by G_t). The new term $h(x_{t+1}^\tau - z_{t+1})$ incorporates uncertainty about the temperature's response to atmospheric CO₂. It has a zero future expectation in the present and is autoregressive with the persistence parameter γ_τ^x . For $\gamma_\tau^x = \eta_\tau = 0$, the stochastic process directly mimics the deterministic radiative forcing relation between CO₂ and generalized (log-) temperature, but what affects the temperature

¹⁴In difference to their applications, the model below makes both autoregression and the shape parameter of the underlying gamma distribution state-dependent, which makes use of the fact that the underlying cumulant generating function is linear not only in last period's state but also in the shape parameter.

¹⁵When combining carbon flow uncertainty with temperature uncertainty also z_t becomes stochastic, but only accounts for the stochastic evolution of carbon, not for the persistent shocks to the temperature response to CO₂ concentrations.

is only the shock's deviation from the expectation $x_{t+1}^\tau - z_{t+1}$. The expected evolution of generalized temperature coincides with that of the deterministic model and, in general, deviations are persistent for $\gamma_\tau^x > 0$. This persistence reflects either the persistence of stochastic feedback processes or the epistemological origin of the uncertainty about the true feedback process. Equation (17) states the variance of the stochastic process, which is proportional to the deviation from the pre-industrial equilibrium $\left(\frac{M_{1,t}}{M_{pre}} - \eta_\tau\right)$ and last period's realization discounted by the persistence factor $(\gamma_\tau^x x_t^\tau)$. In addition, the parameter c scales the variance of the autoregressive gamma process exogenously.

Proposition 5 *Without loss of generality I can set $\delta_\tau = 1$ in the stochastic system of motion (12-18). The optimal carbon tax changes from the deterministic SCC^{det} stated in Proposition 2 to*

$$\begin{aligned} SCC^{unc} &= SCC^{det} \left(1 + \theta_M\right) \left(1 + \frac{\bar{h}}{\sigma_{forc}} (\theta_\tau + \epsilon(c))\right) \\ &\approx SCC^{det} \frac{1 + \frac{\bar{h}}{\sigma_{forc}} \left(\frac{1}{2} \frac{1+\beta\gamma}{1-\beta\gamma} F + \epsilon(c)\right)}{1 - \theta_M^*} \end{aligned} \quad (18)$$

with the uncertainty multipliers

$$\begin{aligned} \theta_M &= \frac{1 - \sqrt{1 - 4\theta_M^*}}{1 + \sqrt{1 - 4\theta_M^*}} \approx \frac{\theta_M^*}{1 - \theta_M^*}. \\ \theta_M^* &= \frac{\alpha\varphi_{M,1}^{det}}{2} \frac{\beta}{M_{pre}} \left[A_M^{M \rightarrow x^2} + A_M^{M \rightarrow \sigma} A_M^{\sigma \rightarrow x^2} + A_M^{M \rightarrow \sigma} \right] \frac{\left([(\mathbf{1} - \beta\Phi)^{-1}]_{1,1} - [(\mathbf{1} - \beta\Phi)^{-1}]_{1,2} \right)^2}{[(\mathbf{1} - \beta\Phi)^{-1}]_{1,1}} \\ &\quad \times \left(1 + \frac{\bar{h}}{\sigma_{forc}} (\theta_\tau + \epsilon(c)) \right) \end{aligned} \quad (19)$$

$$\text{with } A_M^{M \rightarrow x} = \frac{\delta_M^{Mx} \beta}{1 - \gamma_M^x \beta}, \quad A_M^{M \rightarrow \sigma} = \frac{\delta_M^{M\sigma} \beta}{1 - \gamma_M^\sigma \beta}, \quad A_M^{\sigma \rightarrow x} = \frac{\delta_M^{\sigma x} \beta}{1 - \gamma_M^x \beta},$$

and

$$\theta_\tau = \frac{-\log(1 - F(1 + \theta_\tau^*))}{F} - 1 \approx \theta_\tau^* + \frac{1}{2}F \approx \frac{1}{2} \frac{1 + \beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x} F \quad (20)$$

$$\theta_\tau^* = \beta\gamma_\tau^x \frac{1 + F - \sqrt{(1 - F)^2 - 4F \frac{\beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x}}}{1 - F + \sqrt{(1 - F)^2 - 4F \frac{\beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x}}} \approx \frac{\beta\gamma_\tau^x F}{1 - \beta\gamma_\tau^x - F}$$

$$\text{with } F = \alpha\varphi_{\tau,1}^{det} c\bar{h}, \quad \bar{h} = \frac{h}{1 - \beta\gamma_\tau^x},$$

and where $\varphi_{M,1}^{det}$ and $\varphi_{\tau,1}^{det}$ denote the shadow values of atmospheric carbon and generalized temperature under certainty.¹⁶

¹⁶These are the social cost under certainty of carbon and of a (generalized) temperature increase in utils:

Equation (19) states that uncertainty adds two multiplicative factors to the optimal carbon tax. These factors are larger than unity and mutually aggravating. *The first factor* adds the relative increase θ_M , which results from the carbon flow uncertainty. It is a convex function of the uncertainty contribution θ_M^* defined in equation (20). The uncertainty contribution θ_M^* is proportional to the risk aversion α weighted (deterministic) shadow value of atmospheric carbon φ_M^{det} . Note that, under risk *aversion*, both α and φ_M^{det} are negative and the interpretation can neglect these signs as they always come in pairs. First, the risk dependence reveals that it is not merely Arrow-Pratt risk aversion that matters for the social cost of carbon, but how much more Arrow-Pratt risk averse the decision maker is as compared to her desire to smooth consumption over time (α measures this difference, see section 4.2). Second, the relation shows that uncertainty matters more relative to certainty, the worse the climate change problem is in the first place: the deterministic shadow value of atmospheric carbon increases not just the absolute contribution, but also the relative contribution of uncertainty.

At the heart of the uncertainty contribution θ_M^* are the three risk channels abbreviated by $A_M^{M \rightarrow x}$, $A_M^{M \rightarrow \sigma}$, and $A_M^{\sigma \rightarrow x}$, which disentangle the contributions from uncertainty governing conditional expectations and from stochastic volatility. Each of these channels increases in the corresponding δ -parameter that scales the *endogeneity of climate risk*, and with the γ -parameters scaling the *shock persistence*. In particular, I emphasize that only the endogenous uncertainty, which responds to carbon concentrations, affects the optimal policy. The channel $A_M^{M \rightarrow x}$ derives from the conditional expectations process \mathbf{x}_t^M . It is larger, the more an increase in carbon concentration adds to the uncertainty about conditional carbon flow expectations. This $A_M^{M \rightarrow x}$ channel contributes quadratically. The $A_M^{M \rightarrow \sigma}$ channel enters twice. First, it contributes directly through the dependence of stochastic volatility on the carbon concentration. Second, it contributes in combination with $A_M^{\sigma \rightarrow x}$, if the conditional expectations depend themselves on stochastic volatility (if $\delta_M^{\sigma x} > 0$).

Finally, the carbon-flow uncertainty contribution θ_M^* increases (quadratically) in the relative difference between the cost of carbon in the atmosphere (proportional to $[(\mathbf{I} - \beta\Phi)^{-1}]_{1,1}$) and the cost of carbon in the lower ocean & biosphere reservoir (proportional to $[(\mathbf{I} - \beta\Phi)^{-1}]_{1,2}$). The flow uncertainty merely moves carbon between these reservoirs and does not create or annihilate carbon. The cost of a stochastic flow from the ocean into the atmosphere is only harmful to the degree that carbon is more harmful in the atmosphere than in the ocean.

The second uncertainty factor in equation (19) results from the uncertainty of the temperature response to a given evolution of the atmospheric carbon concentration. The two different contributions θ_τ and $\epsilon(c)$ both scale with the size of $\frac{\bar{h}}{\sigma^{forc}}$, the numerator measuring the temperature's sensitivity to the forecast uncertainty, and the denominator measuring the temperature's deterministic response to carbon. The first contribution, θ_τ , captures the

$$\varphi_M^{det} = \frac{\beta \varphi_{\tau,1}^{det} \sigma^{forc}}{M_{pre}} [(\mathbf{I} - \beta\Phi)^{-1}]_{1,1} \text{ and } \varphi_\tau^{det} = -\frac{\xi_0}{1-\beta\kappa} [(\mathbf{I} - \beta\sigma)^{-1}]_{1,1}.$$

direct uncertainty effect resulting from risk and risk aversion. The second contribution, $\epsilon(c)$, results from keeping temperature expectations under uncertainty (approximately) equal to the deterministic temperature evolution calibrated in section 2.5. This second contribution increases in the uncertainty level c . Finding the $\epsilon(c)$ that de-biases the expectations is part of the model calibration.

The approximation of the uncertainty contribution θ_τ in equation (G.18) already conveys the main message how the uncertainty in the temperature-carbon feedback affects the social cost of carbon. It increases and is convex in the discount factor-weighted persistence parameter $\beta\gamma_\tau^x$, and it is approximately proportional in $F = \alpha\varphi_{\tau,1}\frac{c h}{1-\beta\gamma_\tau^x}$. Thus, also the uncertainty governing the temperature response to a given carbon trajectory increases in the risk aversion weighted (deterministic) shadow value, here the social cost of an atmospheric temperature increase. Again, a higher social cost of a temperature increase under certainty implies that a higher *share* of the comprehensive social cost of carbon will be contributed by uncertainty. Moreover, the uncertainty contribution increases in the level of uncertainty c , the sensitivity to forecast uncertainty h , and the discount factor weighted persistence parameter $\beta\gamma_\tau^x$. Evaluating the full expression for the uncertainty contribution shows that the social cost of carbon is not only increasing in F , but convexly increasing in F and its constituents.

5.2 Welfare, Uncertainty, and Learning

This section analyzes the welfare impact of uncertainty and learning. It also introduces a characterization of uncertainty that permits a general analysis of the welfare impact of the various moments of general uncertainty distributions (variance, skewness, kurtosis,...).

5.2.1 Analytic Insights

In the general model, a climate state j is subject to uncertainty and, potentially, to measurement error. The climate state j denotes any of the carbon or temperature layers $M_1, \dots, M_m, \tau_1, \dots, \tau_l$. The measurement error ν_t will be relevant for the case of Bayesian learning and is independently distributed with mean zero and variance $\sigma_{\nu,t}^j$. A second shock ϵ_t^j captures the actual uncertainty we face about the underlying dynamic system. The equations of motion for the climate system are

$$\mathbf{M}_{t+1} = \Phi \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1 + \boldsymbol{\epsilon}_t^M + \boldsymbol{\nu}_t^M \quad (21)$$

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1 + \boldsymbol{\epsilon}_t^\tau + \boldsymbol{\nu}_t^\tau. \quad (22)$$

The random variable ϵ_t^j can be an autoregressive shock process that captures persistent uncertainty, including shocks to higher order moments as in stochastic volatility models.

The random variable ϵ_t^j can also capture epistemological uncertainty and the prior in the case of a Bayesian learning models. I assume that the random variable $\epsilon_t^j = \epsilon_t^j(\kappa_{1,t}^j, \kappa_{2,t}^j, \dots)$ has a distribution that can be characterized uniquely by its cumulant $\kappa_{i,t}^j$ expansion, $i \in \mathbb{N}$. The uncertainty dynamics are pinned down by the cumulants' equations of motion

$$\kappa_{i,t+1}^j = \gamma_i^j \kappa_{i,t}^j + \chi_{i,t}^j, \quad (23)$$

$0 \leq \gamma \leq 1$, for all $i \in \mathbb{N}$, where $\chi_{i,t}^j$ is an sequence of independent random variables (shocks). For example, if only the expected value is non-zero and the shocks are and $\chi_{1,t}^j \sim N(0, \sigma^2)$, then the system characterizes a normally distributed AR(1) shock to state j with persistence γ_1 .

A Bayesian learning model with normally distributed measurement error ν_t (likelihood) and prior ϵ_t gives rise to the updating equation¹⁷

$$\kappa_{1,t+1}^j = \kappa_{1,t}^j + \chi_{1,t}^j \quad \text{with } \chi_{1,t}^j \sim N\left(0, \frac{\sigma_{\epsilon,t}^4}{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2}\right). \quad (24)$$

Writing the updating equation in the form of equation (25) emphasizes the close similarity between learning and a fully persistent AR(1) shock. In addition to equation (25), the variance of ϵ_t does not vanish after observation ($\kappa_{2,t} = \mathbb{V}\text{ar}[\epsilon_t|I_t] > 0$). The variance of this “normal-normal” Bayesian learning model falls *exogenously* over time and does not have to be tracked as an additional state variable (see proof of Proposition 2 for details). However, the shocks $\chi_{1,t}^j$ are correlated with the prevailing epistemological uncertainty and the measurement error, which leads to a different dynamics and welfare loss as compared to the AR(1) model.

Proposition 2 in Appendix H.1 states the general solution for the welfare loss under equations (22-24). Here I focus on two corollaries, the first of which discusses the normally distributed AR and Bayesian learning models.

Corollary 1 *Let uncertainty in equations (22-23) affect state j . A normally distributed first order autoregressive shock process ϵ_t with one step ahead variance σ^2 implies the welfare loss*

$$\Delta W_{normal}^{AR} = \sum_{t=0}^{\infty} \beta^{t+1} \left(\frac{\beta}{1-\gamma\beta}\right)^2 \alpha \varphi_j^2 \frac{\sigma_x^2}{2} = \frac{\beta}{1-\beta} \left(\frac{\beta}{1-\gamma\beta}\right)^2 \alpha \varphi_j^2 \frac{\sigma^2}{2}. \quad (25)$$

A Bayesian learning model with normally distributed prior $\epsilon_t \sim N(\mu_{\epsilon,t}, \sigma_{\epsilon,t}^2)$ and measurement

¹⁷The standard way of writing the Bayesian updating equation for the mean is $\mu_{\epsilon,t+1} = \frac{\sigma_{\epsilon,t}^2}{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2} \mu_{\epsilon,t} + \frac{\sigma_{\nu,t}^2}{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2} z_t$ with observation $z \sim N(\mu_{\epsilon,t}, \sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2)$. Defining $\chi_{1,t}^j = \frac{\sigma_{\nu,t}^2}{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2} (z_t - \mu_{\epsilon,t})$ and writing $\kappa_{1,t}$ for the mean $\mu_{\epsilon,t}$ delivers equation (25). Note that the observational variable z is defined in equations (22-23). For example, in the case of uncertain atmospheric carbon content, the observation z is $M_{t+1} - \Phi_{1,\cdot} \mathbf{M}_t - (\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo})$.

error $\nu_t \sim N(0, \sigma_{\nu,t}^2)$ implies the welfare loss

$$\Delta W^{Bayes} = \sum_{t=0}^{\infty} \beta^{t+1} \left(\frac{1}{1-\beta} \right)^2 \Omega_t^2 \alpha \varphi_j^2 \frac{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2}{2} \quad (26)$$

with $\Omega_t \equiv \frac{\sigma_{\epsilon,t}^2}{\sigma_{\nu,t}^2 + \sigma_{\epsilon,t}^2} + (1-\beta) \frac{\sigma_{\nu,t}^2}{\sigma_{\nu,t}^2 + \sigma_{\epsilon,t}^2}$.

If uncertainty affects more than one state, then the welfare loss is additive.

In both models, the welfare loss is proportional to (intrinsic) risk aversion α and the square of the state's shadow value, which is the same as under certainty. This shadow value is $\varphi_j = \varphi_\tau$ if the model characterizes the uncertainty about atmospheric temperature feedbacks. The model requires two shocks to model the uncertainty governing the carbon flow between the atmosphere and the adjacent sinks because carbon only moves from one reservoir to the other. Here, a shock to the atmospheric carbon is accompanied by a perfectly negatively correlated shock to the adjacent layer, giving rise to the shadow value is $\varphi_j = \varphi_{M,1} - \varphi_{M,2}$ (see proof of Corollary 1 for details).

The welfare loss is proportional to the variance. In the AR model, this variance is constant by assumption. Then, the infinite sum over future contributions can be summarized by the factor $(1 - \beta)^{-1}$. This factor yields a similar time sensitivity as observed for the carbon tax in Proposition 2. However, the welfare loss from the AR uncertainty exhibits an additional factor $(1 - \gamma\beta)^{-2}$ that makes it even more sensitive to the persistence weighted time preference: the more persistent the shock, the longer its implications, and the more relevant is time preference in determining the welfare loss.

The Bayesian learning model has the time varying variance $\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2$ combining the prior's and the likelihood function's uncertainty. The prior's variance declines as the decision maker improves her estimate of the true system dynamics. The most striking difference is that the Bayesian model gives rise to the additional time sensitivity factor $(1 - \beta)^{-1}$ that corresponds to an AR model with unit persistence. Intuitively, a learning shock updates the Bayesian future for good. Yet, the "AR" shocks in the Bayesian model are correlated with the epistemological uncertainty and fall over time, which gives rise to the additional factor Ω^2 . The factor Ω is a weighted mean of unity and $1 - \beta$. Initially, when the prior is still vague ($\sigma_{\epsilon,t}$ large), the factor is close to unity and the Bayesian decision maker updates her long-run belief. If she is patient, this long-run update moves her welfare by much. Once the decision maker has a better picture about the true system dynamics, the term Ω increasingly cancels the additional time sensitivity. Once, epistemological uncertainty is gone ($\sigma_{\epsilon,t} \ll \sigma_{\nu,t}$), all that remains is an iid measurement error, and Ω cancels the additional sensitivity to time preference.

Kelly & Kolstad (1999) and Karp & Zhang (2006) employ the simple Bayesian learning model for the assessment of climate change feedbacks and damages. Corollary 1 extends immediately to general Kalman filter models. This extension is empirically relevant to capture

a more realistic learning and temperature adjustment speed (Jensen & Traeger 2013). The second corollary based on the general system of equations (22-24) focuses on the welfare response to higher order moments. Zooming in on these characteristics, the corollary assumes that there is no correlation between the prevailing epistemological uncertainty in period t and the period t shock and ignores measurement error.

Corollary 2 *Let equations (22-24) characterize uncertainty. Let $\nu_t = 0$ for all t and let the shocks to the cumulants $\chi_{i,t}^j$ be independent of the prevailing epistemological uncertainty ϵ_t^j .*

$$\Delta W^{unce} = \underbrace{\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \kappa_{i,0}^j}_{\text{Epistemological}} + \sum_{t=0}^{\infty} \frac{\beta^{t+1}}{\alpha} \left[\underbrace{\sum_{i=1}^{\infty} G_{\chi_i^j}(\alpha \varphi_{\kappa,i}^j)}_{\text{Future Shocks}} \right] \quad (27)$$

where the shadow cost of the i^{th} cumulant is

$$\varphi_{\kappa,i}^j = \frac{\beta}{1 - \beta \gamma_i^j} \frac{(\alpha \varphi_j)^i}{i! \alpha}, \quad (28)$$

and φ_j denotes the shadow value of the state subject to the uncertainty. If more than one state j are uncertain, then the welfare loss is additive over j .

The first component of the welfare loss is an ‘‘epistemological contribution’’. It sums the shadow value weighted cumulants of the prevailing uncertainty distribution in the present ($t = 0$). The second contribution derives from the future shocks.

The *general vector autoregressive shock model* corresponds to the case where all but the first cumulant (expected value) are zero: $\mathbb{E}[\chi_{i,t}^j | I_t] = 0$ for all t and $i > 1$. The decision maker knows the exact state of the dynamic system at the beginning of every period. Assuming that shocks $\chi_{1,t}^j$ are independently and identically distributed translates the infinite sum into a factor $\frac{1}{1-\beta}$ and delivers the welfare loss¹⁸

$$\Delta W_{general}^{AR} = \frac{\beta}{\alpha(1-\beta)} G_{\chi}(\alpha \varphi_{\kappa,1}^j) = \frac{\beta}{\alpha(1-\beta)} \sum_{l=1}^{\infty} \kappa_l^{\chi} \frac{(\alpha \varphi_{\kappa,1}^j)^l}{l!}, \quad (29)$$

where the random variable χ has the generic distribution of an individual shock in the sequence $\chi_{1,t}^j$ and the cumulants κ_l^{χ} (not $\kappa_{i,t}^j$) characterize the moments of the *shock*. These autoregressive shocks are to the mean, and changes in the mean are valued at $\varphi_{\kappa,1}^j = \frac{\beta}{1-\beta\gamma_1^j} \varphi_j$, where γ_1^j is the autoregressive coefficient. The variance of the shock to the mean reduces welfare proportional to $\alpha \varphi_{\kappa,1}^j{}^2 = \left(\frac{\beta}{1-\beta\gamma_1^j}\right)^2 \varphi_j^2$, as derived in Corollary 1. All estimates of the climate sensitivity distribution exhibit both skewness and kurtosis. Equation (30) shows that the skewness κ_3^{χ} of the shock delivers a welfare loss proportional to $\alpha^2 \left(\frac{\beta}{1-\beta\gamma_1^j}\right)^3 \varphi_j^3$, making it

¹⁸By the assumption that expected one step ahead dynamics in the present coincides with the deterministic model (equations ??-??) it is $\mu_{\epsilon} = \kappa_{1,0} = 0$.

even more sensitive to (persistence weighted) time preference, and perhaps non-surprisingly to risk aversion. Kurtosis, a measure of “heaviness in the tails” adds yet another power to each of these sensitivities.

We cannot predict with certainty today how uncertain we will be about carbon and temperature dynamics a decade or a century from today. Thus, the variance of our uncertainties is itself a stochastic process. A *stochastic volatility* model adds shocks to the second cumulant, making the variance of the uncertainty distribution itself a stochastic process. By Corollary 2, changes of the variance are evaluated at the shadow value $\varphi_{\kappa,2}^j = \frac{\beta}{1-\beta\gamma_2^j} \frac{\alpha\varphi_j^2}{2}$. In contrast to the shadow value of a change in the mean, a change of the variance is evaluated proportional to risk aversion, and proportional to the square of the state’s shadow value. Independently and identically distributed shocks $\chi_{2,t}^j$ to the variance lead to a welfare loss of $\frac{\beta}{\alpha(1-\beta)} G_{\chi_2^j}(\alpha\varphi_{\kappa,2}^j)$. The variance of such a shock contributes a welfare loss proportional to $\left(\frac{\beta}{1-\beta\gamma_2^j}\right)^2 \alpha^3 \varphi_j^4$, and skewness has an impact proportional to $\left(\frac{\beta}{1-\beta\gamma_3^j}\right)^3 \alpha^5 \varphi_j^6$. As observed for the shocks to the mean, each order of the shock’s cumulant adds a power to the sensitivity to (persistence weighted) time preference. However, for shocks to the variance, each order of the shock’s cumulant adds *two* powers to the sensitivity to risk aversion and to the shadow value of the uncertain state.

The *purely epistemological contribution* $\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \kappa_{i,0}^j$ in equation (28) adds a loss that prevails even in the absence of future shocks. By assumption, today’s mean vanishes (equations ??-??). Today’s uncertainty about the variance implies a welfare loss of $\varphi_{\kappa,2}^j \kappa_{2,0}^j = \frac{\beta}{1-\beta\gamma_2^j} \frac{\alpha\varphi_j^2}{2} \kappa_{2,0}^j$, and the skewness of today’s belief about the future implies a loss of $\varphi_{\kappa,3}^j \kappa_{3,0}^j = \frac{\beta}{1-\beta\gamma_3^j} \frac{\alpha^2\varphi_j^3}{6} \kappa_{3,0}^j$. The sensitivity to time preference contrasts with that of the future shock contributions. Here, the sensitivity to persistence-weighted time preference does *not* increase in the moments of the distribution; “fat tails” (kurtosis and higher moments) in today’s epistemological distribution are not more sensitive to time preference than are the mean contribution or the contribution of the variance.

In some models, the *persistence of epistemological uncertainty* can be lower for a fat tail of the distribution, see Kelly & Tan (2015b) for an application to climate sensitivity. The perhaps simplest stylized model of epistemological uncertainty without future shocks assumes an exponential decay of the uncertainty distribution. Because the i^{th} cumulant is homothetic of degree i , an exponential uncertainty decay that leads to ϵ_{t+1} distributed as $\gamma\epsilon_t$ implies that $\gamma_{i,t}^j = (\gamma^j)^i$. Then, the sensitivity to time preference of higher moments is weighted down by the falling factor $(\gamma^j)^i$. The faster decay of skewness, kurtosis, and higher moments implies a lower sensitivity to time preference.

In general, the l^{th} cumulant of a shock to the j^{th} cumulant of the uncertainty distribution causes a welfare loss proportional to l^{th} power of the (persistence weighted) time preference, to the $l \cdot i^{th}$ power of the temperature’s shadow value, and to the $l \cdot i - 1^{th}$ power of risk aversion. In summary, time sensitivity grows with the order of the shock’s moment, whereas

sensitivity to risk aversion and the state’s deterministic shadow value grow as the product of the orders of the shock’s and shocked moment. Finally, it is not risk aversion in the Arrow Pratt sense that drives the welfare loss. Instead, the intrinsic risk aversion measure $-\alpha$ characterizes by how much Arrow Pratt risk aversion dominates the agent’s aversion that results from her mere desire to smooth consumption over time.¹⁹

5.2.2 Quantification

I use a combination of 20 estimates of climate sensitivity from Meinshausen et al. (2009) to calibrate *temperature uncertainty* for a given carbon trajectory. These estimates derive from different research groups and use a variety of methodological approaches. Lacking a direct estimate of the probabilities, I use a study subjecting 18 different carbon cycle models to different shocks to calibrate uncertainty about the emissions flow uncertainty between the atmosphere and the ocean-biosphere sink.²⁰ The estimates of the welfare loss assume (a stochastic version of) the business as usual scenario of the DICE 2013 model. The estimates of the optimal carbon tax assume option policy.²¹ I evaluate all carbon cycle uncertainty scenarios with a risk aversion coefficient of $\alpha = -1.5$ (see section 4.2 and Appendix F.1).

Absent better information, a normal distribution seems reasonable to characterize the uncertain dynamics of the *carbon flow* between the atmosphere and the ocean-biosphere sinks, and I use Corollary 26 to evaluate its welfare impact. Joos et al.’s (2013) simulations suggests a high persistence of $\gamma = 0.997$ in equation (26) to model future carbon flow as an AR(1) process. Based on these simulations, Appendix H.4 argues for a variance of the decadal shock to the carbon exchange between atmosphere and ocean-biosphere of approximately $\sigma_\chi = 20$ Gt per decade. Figure 4, left panel, shows the resulting uncertainty governing the atmospheric carbon concentration along the business as usual scenario of DICE 2013. For epistemological uncertainty, I assign this variance to the Bayesian prior ($\sigma_{\epsilon,0} = 20$ Gt per decade), and set the “measurement error” to $\sigma_\nu = 10$ Gt (corresponding approximately to the currently “missing” carbon flow). This combination of prior and measurement error implies a remaining epistemological uncertainty with a standard deviation of 4.4Gt per decade after 50 years and of 2.6Gt per decade after 150 years. The uncertainty results in an approximate present value welfare loss of 110 billion USD in the AR(1) model, and 30 billion USD in the Bayesian model. For a comparison of magnitude, the annual NASA budget is about 20 billion USD.

¹⁹Observe that shadow values are negative and the product $\alpha^{l-i-1} \varphi^{j^{l-i}}$ always implies a welfare loss for a positive degree of intrinsic aversion $-\alpha > 0$.

²⁰Note that the carbon cycle exhibits predicted uncertain feedbacks that can hardly be observed at present concentration and temperature levels. Thus, I rely on a large set of model output rather than a simplistic evaluation of historical data. In particular for the carbon cycle, any estimate of uncertainty is itself uncertain and subject to model bias.

²¹The welfare loss in the optimal regime is not specified without detailed assumptions about the energy sector (equation 2) because it depends on the sectors’ responses to the tax.

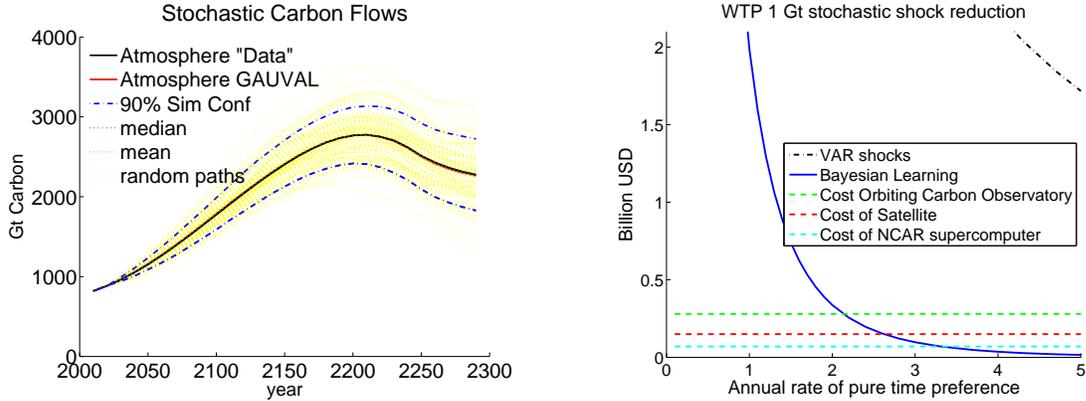


Figure 4: The graph on the left shows the evolution of atmospheric carbon for the DICE 2013 business as usual emission scenario. Decadal shocks with a standard deviation of $\sigma_\chi = 20$ Gt per decade change the flow between the atmosphere and the carbon sinks with a persistence of $\gamma_M = 0.997$ that is calibrated to the carbon cycle comparison study by Joos et al. (2013). The deterministic DICE evolution (5 year time steps, “Data”), the deterministic ACE evolution (10 year time steps), and the mean and the median of 1000 uncertain trajectories are hardly distinguishable. The right graph depicts the willingness to pay for a 1Gt uncertainty reduction. In the Bayesian learning case, the reduction is in the measurement error, increasing the speed of learning. In the case of the vector autoregressive shock model (“VAR”), the willingness to pay is based on a physical reduction of carbon flow stochasticity (e.g., a co-benefit of emission reductions).

Figure 4, right panel, states the willingness to pay for a 1Gt reduction of the decadal standard deviation as a function of pure time preference. In the case of the VAR model, the uncertainty reduction lowers the physical stochasticity of the carbon flows ($\sigma_\chi = 10\text{Gt} \rightarrow 9\text{Gt}$).²² In the case of the Bayesian model, the uncertainty reduction lowers the measurement error and increases the speed of learning ($\sigma_\nu = 10\text{Gt} \rightarrow 9\text{Gt}$). The figure compares the welfare gain from better measurement and faster learning to the costs of a common satellite (~ 150 million USD), NASA’s Orbiting Carbon Observatory (~ 280 million), and the National Center for Atmospheric Research’s recent supercomputer (~ 70 million). For the standard calibration of the time preference these investments are worth the (global) welfare gain. For an annual rate of time preference around 3% even the global welfare gain might no longer outweigh their costs in ACE.²³ A normatively or long-run risk founded rate of $\rho = 0.1\%$ increases the willingness to pay for a decadal 1Gt stochasticity reduction to approximately

²²Reducing the carbon flow’s decadal standard deviation σ_χ by 1 Gt reduces the welfare loss by the fraction $\frac{2}{\sigma_\chi} + \frac{1}{\sigma_\chi^2}$. This formula solves $x = \frac{\Delta W_{\sigma_\chi}^{VAR,normal} - \Delta W_{\sigma_\chi - 1}^{VAR,normal}}{\Delta W_{\sigma_\chi}^{VAR,normal}}$ for x . The graph is only visible in the upper right corner: the payoff of the physical stochastic shock reduction is more valuable than a reduction of measurement error that accelerates learning.

²³This value is most sensitive to the initial prior. Lowering initial uncertainty to $\sigma_{\epsilon,0} = 10\text{Gt}$ lowers the welfare loss to approximately 10 billion USD. To obtain the same 110 billion USD welfare loss as from the VAR setting, I would have to raise initial uncertainty of the prior to as much as a 40Gt standard deviation. Note that the NASA’s Orbiting Carbon Observatory is the investment closest to a direct reduction of measurement error to improve learning. Slowly coming out of its calibration phase, the ultimate precision is still unclear.

110 billion USD. In fact, for such a low rate of pure time preference the uncertainty induced total present value welfare loss increase to 60 trillion USD, or 73% of world output, in the case of Bayesian learning. In contrast, the welfare loss produced by the AR(1) process under the low discount rate rises only to 2 trillion USD, showing the higher initial sensitivity of the Bayesian learning to pure time preference derived in Corollary 26.

I conclude that the absolute welfare costs from uncertainty over the carbon flows are small to moderate compared to the deterministic contributions discussed in section 3. A higher carbon shock implies both a higher temperature, leading to more convex damages, but also a higher satiation of the CO₂'s absorption spectrum, leading to a lower marginal impact of the last unit of emissions. These two effects largely offset each other. Under the standard discounting calibration, this welfare cost is 2-3 orders of magnitude lower than the welfare loss from present CO₂ concentrations (and the resulting warming that is already in the pipeline). Given the higher sensitivity to time preference, the welfare loss from uncertainty hesitantly catches up to a similar order of magnitude for the patient decision maker.²⁴

ACE's equations of motion are exponential in temperature (linear in $\tau_1 = \exp(\xi_1 T_1)$). Thus, even a normal distribution of temperature can translate into a log-normal distribution in the linear equations of motion. By Corollary 2, the resulting welfare loss is proportional to the cumulant generating function. The cumulant generating function of the log-normal distribution is infinite. Hence, I can easily set up a model that delivers an infinite welfare loss from climate sensitivity uncertainty. This result takes Weitzman's (2009a) "dismal theorem" and Millner's (2013) extension from their stylized frameworks into a full-fledged and well-calibrated integrated assessment model. Here, even the thin-tailed normal distribution as opposed to the fat-tailed prior in Weitzman (2009a) can blow up the welfare loss through its translation into economic damages.²⁵ In a stochastic numeric implementation of DICE, Kelly & Tan (2015b) show that fat tails on climate sensitivity are accompanied by relatively quick learning (see also intuition at end of section 5.2). Such quick learning would make fat tails less relevant (though actual learning seemed rather slow so far). The present result applies even in the case of thin tails. I interpret these result as a mere call of attention to the analysis of uncertainty in climate change and for a call of careful guesstimates of what life and welfare might be like for global warming above 5, 10, or even 15C. Unfortunately, these scenarios are highly relevant for today's climate change evaluation and policy and yet impossible to derive from present data.

The present quantitative illustration of temperature uncertainty will follow the hum-

²⁴The 14 trillion USD loss from the present atmospheric CO₂ concentration increases to 160 trillion USD for $\rho = 0.1\%$, which is at least the same order of magnitude as the 60 trillion USD loss from Bayesian uncertainty.

²⁵Figure 2 shows that for temperature increases up to 12C, ACE's base case damage specification delivers damages lower than DICE. More than that, the "dismal result" holds for any $\xi_0 > 0$, implying that I can make damages at any given temperature level arbitrarily small and still find an infinite welfare loss from temperature uncertainty.

ble (or possibly crazy) attempt to obtain an estimate of the welfare loss from uncertainty limiting global temperature rise to at most 10C. This range corresponds to the domain of the probabilistic climate sensitivity estimates, and to damages of up to some 20% of world output. The resulting welfare loss is a lower bound.

I calibrate the shock uncertainty to produce the climate sensitivity distribution in the infinitely long run. I split the uncertainty into an epistemological fraction ζ and a shock-based long-run fraction $1 - \zeta$. At any given point in time the actual forecast uncertainty will then be lower than the climate sensitivity distribution because epistemological uncertainty falls over time and the shocks only build up the fraction $1 - \zeta$ in the long-run. Third, I omit possible contributions from stochastic volatility or shocks to higher order cumulants. Fourth, I pick $\alpha = -1$ at the lower end of measured risk aversion.

My “baseline” scenario assumes a persistence $\gamma = 0.9$ of both epistemological uncertainty and shocks to the mean, an equal split of overall climate sensitivity uncertainty between the epistemological and the shock contributions ($\zeta = \frac{1}{2}$), and the standard discount rate calibration to IMF 2015 data ($\rho = 1.75\%$). These assumption result in an overall welfare loss from climate sensitivity uncertainty of 16 trillion USD, approximately one year of US output. Initial epistemological uncertainty and the stochastic shocks contribute almost equal shares to this loss. As a consequence, attributing a larger or smaller share of the uncertainty to shocks and future updating hardly changes the welfare loss.

Varying the persistence of shocks and of epistemological uncertainty between a lower value of $\gamma = 0.7$ and the higher value $\gamma = 0.997$ calibrated for the carbon flow uncertainty varies the welfare loss between 11.5 and 20 trillion USD. A reduction of pure time preference to $\rho = 0.1\%$ in the “baseline” scenario increases the loss to over 700 trillion USD or 8.5 years of world output. This factor 40 increase is significantly larger than the response of the carbon tax to the change in pure time preference. The scenario also confirms the theoretical finding that the future shock contributions are more sensitive to time preference than mere epistemological uncertainty: 95% of this welfare loss derives from the future shocks. I reiterate that the future shock component of the present model reflects the learning shocks that make the Bayesian model particularly sensitivity to pure time preference. Thus, the finding fleshes out that the high sensitivity to pure time preference in the epistemological models does not result from the mere presence of epistemological uncertainty, but from its anticipated updating and the corresponding long-run welfare impact of the shocks.

Finally, the lower bound of the welfare loss from uncertainty over the climate’s sensitivity is 2-3 orders of magnitude higher than the best guess of the welfare loss from uncertainty over the carbon flows. A clear quantitative message from economics to science is to shift more attention to the feedback processes on the temperature side.

6 Conclusions

ACE is an integrated assessment model of climate change that matches scientific climate models just as well as do numeric models used in policy advising. It derives the optimal carbon tax and welfare loss in closed form. ACE merges Golosov et al.’s (2014) framework with a standard carbon cycle, radiative forcing, temperature dynamics, risk attitude, and different uncertainty frameworks. The resulting model closely resembles (a stochastic version) of the widely used integrated assessment model DICE. It improves the market calibration by disentangling intertemporal substitution and the low risk-free discount rate from risk aversion and risk premia.

The deterministic model finds a market-based optimal carbon tax of 57 USD per ton of carbon (15 USD per ton of CO₂), using a standard calibration approach. The closed-form solution shows that the carbon cycle’s persistence is the main multiplier of the SCC (almost a factor 4), whereas temperature dynamics cause a reduction of the optimal tax (by 40%). Analyzing the system’s shadow values, ACE shows that the welfare loss from the present increase in carbon concentrations is higher than from the present increase in global temperature (independent of future policy). Like Golosov et al.’s (2014) model, ACE implies a flat marginal benefit curve from mitigation. The finding underpins the advantages of a carbon tax over a cap and trade mechanism to regulate the climate externality. Another convenient consequence is that we do not have to know the highly complex mitigation technology frontier for optimal carbon regulation. Finally, optimal mitigation effort is independent of whether we followed business as usual or optimal policy in the past. If we “sinned” in the past, the optimal policy will not tell us to repent, but to live with the (perpetually) persisting consequences in the future.

A wide-spread belief is that the optimal carbon tax is sensitive to the overall consumption discount rate, but not to its individual constituents. In contrast, I prove in the present well-calibrated setting that mass conservation in the carbon cycle makes the optimal carbon tax highly sensitive to the rate of pure time preference ($\approx \frac{1}{\rho}$), whereas proportionality of damages to output make it insensitive to growth related discounting. The sensitivity to pure time preference weighs particularly strong because recent asset pricing approaches and overlapping generations based calibration formulas suggest much lower rates of pure time preference than the 1.75% calibrated here following a standard approach. These more sophisticated approaches support rates as low as the 0.1%, which the Stern (2007) Review used for normative reasons. Such a pure rate of time preference increases the optimal carbon tax tenfold, with a sevenfold increase resulting from the carbon cycle interaction.

I employ ACE to advance our understanding of how current climate policy should respond to future climate uncertainty. I show that the “risk premia” are positive for both carbon flow and temperature uncertainty. Both uncertainty contributions to the optimal carbon tax are mutually aggravating, implying that joint uncertainty has a stronger impact

on optimal policy than the sum of each individual channel. Moreover, the uncertainty factor multiplying the deterministically optimal tax increases convexly with risk aversion, the deterministic shadow value of atmospheric carbon and atmospheric temperature, the persistence of uncertainty, the discount factor, and “uncertainty” itself. Moreover, policy is mostly affected by endogenous uncertainty rather than by exogenous shocks.

A preliminary quantification, for which I have to refer to the official working paper version to be released soon, shows the following. Uncertainty about conditional expectations contributes much more than stochastic volatility. This finding reflects that carbon flow irregularities at the decadal scale are less relevant than long-term uncertainties in carbon accumulation. Overall, the impact of carbon flow uncertainty on mitigation policy is small, and much smaller than the impact of temperature uncertainty. Uncertainty about the warming resulting from a given carbon trajectory results in an approximate increase of the optimal carbon tax by 20% under the standard calibration. Reducing pure time preference from 1.75% to 1%, as suggested in a recent expert elicitation, increases the SCC increase under uncertainty from 20% to approximately 360%. In contrast to suggestions that uncertainty is more relevant for the optimal carbon tax than discounting, I find that the uncertainty contribution to the optimal carbon tax is even more sensitive to the discount rate than its impact on the deterministic contribution.

The paper analyses the welfare impact of stochasticity, uncertainty, and learning in a model that allows for general uncertainty distributions. The moments (cumulants) of either of the uncertainty distributions reduce welfare proportional to the corresponding powers of the risk aversion weighted shadow value of a change in warming or carbon flows. The applicable measure of risk aversion is not the Arrow-Pratt measure, but a measure of intertemporal or intrinsic risk aversion. This measure characterizes how much more averse a decision maker is to risk than to deterministic consumption fluctuations. The welfare loss’ sensitivity to pure time preference increases for higher moments of the shock’s distribution, but slightly decreases for higher moments of the present epistemological uncertainty that captures the decision maker’s lack of knowledge. Bayesian updates to the epistemological uncertainty act initially like fully persistent shocks as they carry information on a permanent change of the system dynamics. Therefore, the overall welfare loss from uncertainty under anticipated learning is highly sensitive to pure time preference.

In the standard calibration, the welfare loss from carbon cycle uncertainty is in the order of a hundred billion USD. The willingness to pay for a reduction of the measurement error and for accelerated learning is in the order of half a billion per Gt C of decadal resolution (the cost of a few satellites and a supercomputer). The Bayesian model’s sensitivity to pure time preference increases this willingness to pay to 110 billion USD for a rate of pure time preference of 0.1%, which is more than the value of full uncertainty elimination in the standard calibration. Uncertainty about the temperature response to a given CO₂ path causes a lower bound welfare loss that is 2 to 3 orders of magnitudes larger, about one year

of US output in the standard calibration, and over 8 years of world output under a rate of pure time preference of 0.1%.

Governments and research institutions are spending large amounts to better understand carbon flows. An immediate conclusion is that better assessments of the temperature feedback response have a significantly higher social payoff. The intuition is the following. Every additional ton of carbon in the atmosphere traps less energy than the preceding ton. This “decreasing harmfulness” of CO₂ to temperature offsets the convexity of damages from the implied warming and the decreasing marginal utility (governing intertemporal trade-offs). Thus, negative and positive shocks in the carbon flow would offset each other if it was not for (disentangled) risk aversion. Risk aversion implies a moderate willingness to pay for a risk reduction. In contrast, temperature feedbacks operate directly on temperatures. Because of the convex damage function, high temperature realizations cause more loss than is gained back from low realizations. In expectation, the shocks reduce overall welfare, an effect that is only amplified by risk aversion.

The present paper paves the way for a wide array of analytic and quantitative research. ACE can be generalized for regional analysis, to examine adaptation, to analyze detailed damage channels like ocean-acidification or sea level rise, and to evaluate benefits from climate engineering projects. The present paper specifies the optimal carbon tax for a large class of energy sectors. Specifying their details permits an easy analysis of the sectoral emission response to optimal policy under technological uncertainty. Climate change is an intergenerational problem. The present paper focuses on market-based evaluation, following common practice of policy advising in the US. ACE also lends itself to a normatively motivated analysis. ACE’s major virtue is to combine quantitative analysis with analytic insight. Any analytic approach has its limitations in the non-linearities and interactions it can handle. The model serves best as a benchmark, guiding and helping to enlighten fine-tuned quantitative numeric research.

References

- Anderson, E., Brock, W., Hansen, L. P. & Sanstad, A. H. (2014), ‘Robust analytical and computational explorations of coupled economic-climate models with carbon-climate response’, *RDCEP Working Paper No.13-05* .
- Anthoff, D. & Tol, R. S. (2014), The climate framework for uncertainty, negotiation, and distribution (fund), technical description, version 3.9, Technical report, <http://www.fund-model.org/>.
- Bansal, R., Kiku, D. & Yaron, A. (2010), ‘Long-run risks, the macro-economy and asset prices’, *American Economic Review: Papers & Proceedings* **100**, 542–546.

- Bansal, R., Kiku, D. & Yaron, A. (2012), ‘An empirical evaluation of the long-run risks model for asset prices’, *Critical Finance Review* **1**, 183–221.
- Bansal, R. & Yaron, A. (2004), ‘Risks for the long run: A potential resolution of asset pricing puzzles’, *The Journal of Finance* **59**(4), 1481–509.
- Brock, W. A. & Mirman, L. J. (1972), ‘Optimal economic growth and uncertainty: The discounted case’, *Journal of Economic Theory* **4**, 479–513.
- Calel, R. & Stainforth, D. A. (2015), ‘Transparency in integrated assessment modeling’, *Working Paper* .
- Chen, X., Favilukis, J. & Ludvigson, S. C. (2013), ‘An estimation of economic models with recursive preferences’, *Quantitative Economics* **4**, 39–83.
- Epstein, L. G. & Zin, S. E. (1991), ‘Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis’, *Journal of Political Economy* **99**(2), 263–86.
- Gerlagh, R. & Liski, M. (2012), Carbon prices for the next thousand years, Cesifo working paper series, CESifo Group Munich.
- Gerlagh, R. & Liski, M. (2014), Carbon prices for the next hundred years, Cesifo working paper series, CESifo Group Munich.
- Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2014), ‘Optimal taxes on fossil fuel in general equilibrium’, *Econometrica* **82**(1), 41–88.
- Gourieroux, C. & Jasiak, J. (2006), ‘Autoregressive gamma processes’, *Journal of Forecasting* **25**, 129–152.
- Hassler, J. & Krusell, P. (2012), ‘Economics And Climate Change: Integrated Assessment In A Multi-Region World’, *Journal of the European Economic Association* **10**(5), 974–1000.
- Heal, G. (1984), Interactions between economy and climate: A framework for policy design under uncertainty, *in* V. K. Smith & A. D. White, eds, ‘Advances in Applied Microeconomics’, Vol. 3, JAI Press, Greenwich, CT, pp. 151–168.
- Hepburn, C. (2006), Discounting climate change damages: Working notes for the Stern review, Working note.
- 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.
- IPCC (2013), *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Iverson, T. (2013), ‘Optimal carbon taxes with non-constant time preference’, *MPRA Paper* **49588**.
- Jensen, S. & Traeger, C. (2013), ‘Optimally climate sensitive policy: A comprehensive evaluation of uncertainty & learning’, *Department of Agricultural and Resource Economics, UC Berkeley* .
- Joos, F., Roth, R. & Weaver, A. J. (2013), ‘Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis’, *Atmos. Chem. Phys.* **13**, 2793–2825.
- Karp, L. (2013), ‘Provision of a public good with altruistic overlapping generations and many tribes’, *Working Paper* .
- Karp, L. & Rezai, A. (2017), ‘Asset prices and climate policy’, *Working Paper* .
- Karp, L. & Zhang, J. (2006), ‘Regulation with anticipated learning about environmental damages’, *Journal of Environmental Economics and Management* **51**, 259–279.
- Karp, L. & Zhang, J. (2012), ‘Taxes versus quantities for a stock pollutant with endogenous abatement costs and asymmetric information’, *Economic Theory* **49**, 371–409.
- Kelly, D. L. & Kolstad, C. D. (1999), ‘Bayesian learning, growth, and pollution’, *Journal of Economic Dynamics and Control* **23**, 491–518.
- Kelly, D. L. & Tan, Z. (2015a), ‘Learning and climate feedbacks: Optimal climate insurance and fat tails’, *Journal of Environmental Economics and Management* **72**, 98–122.
- Kelly, D. L. & Tan, Z. (2015b), ‘Learning and climate feedbacks: Optimal climate insurance and fat tails’, *University of Miami Working paper* .
- Le, A., Singleton, K. J. & Dai, Q. (2010), ‘Discrete-time affine term structure models with generalized market prices of risk’, *Review of Financial Studies* **23**, 2184–2227.

- Leach, A. J. (2007), ‘The climate change learning curve’, *Journal of Economic Dynamics and Control* **31**, 1728–1752.
- Li, X., Narajabad, B. & Temzelides, T. (2014), ‘Robust dynamic optimal taxation and environmental externalities’, *CESifo Working Paper Series* **4562**.
- Meinshausen, M., Meinshausen, N., Hare, W., Raper, S., Frieler, K., Knutti, R., Frame, D. & Allen, M. (2009), ‘Greenhouse-gas emission targets for limiting global warming to 2c’, *Nature* **458**, 1158–1162.
- Meinshausen, M., Raper, S. & Wigley, T. (2011), ‘Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, magicc6 - part 1: Model description and calibration’, *Atmospheric Chemistry and Physics* **11**, 1417–1456.
- Millner, A. (2013), ‘On welfare frameworks and catastrophic climate risks’, *Journal of Environmental Economics and Management* **65**(2), 310–325.
- Nakamura, E., Steinsson, J., Barro, R. & Ursua, J. (2013), ‘Crises and recoveries in an empirical model of consumption disasters’, *American Economic Journal: Macroeconomics* **5**(3), 35–74.
- Newell, R. G. & Pizer, W. A. (2003), ‘Regulating stock externalities under uncertainty’, *Journal of Environmental Economics and Management* **45**, 416–432.
- 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.
- Pindyck, R. S. (2013), ‘Climate change policy: What do the models tell us?’, *Journal of Economic Literature* **51**(3), 860–872.
- Rezai, A. & der Ploeg, F. V. (2016), ‘Intergenerational inequality aversion, growth, and the role of damages: Occams rule for the global carbon tax’, *Journal of the Association of Environmental and Resource Economists* **3**(2), 493522.
- Roe, G. H. & Bauman, Y. (2013), ‘Climate sensitivity - should the climate tail wag the policy dog’, *Gerard H. Roe and Yoram Bauman* **117**, 647–662.
- Schneider, M., Traeger, C. P. & Winkler, R. (2013), ‘Trading off generations: Infinitely lived agent versus oig’, *European Economic Review* **56**, 1621–1644.

- Stern, N., ed. (2007), *The Economics of Climate Change: The Stern Review*, Cambridge University Press, Cambridge.
- Traeger, C. (2012a), ‘Once upon a time preference - how rationality and risk aversion change the rationale for discounting’, *CESifo Working Paper* **3793**.
- Traeger, C. P. (2012b), ‘A 4-stated dice: Quantitatively addressing uncertainty effects in climate change’, *Environmental and Resource Economics* **59**, 1–37.
- Traeger, C. P. (2014), ‘Capturing intrinsic risk aversion’, *Working Paper* .
- van den Bijgaart, I., Gerlagh, R. & Liski, M. (2016), ‘A simple formula for the social cost of carbon’, *Journal of Environmental Economics and Management* **77**, 75–94.
- 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. (1974), ‘Prices versus quantities’, *Review of Economic Studies* **41**(4), 477–91.
- Weitzman, M. (2009a), ‘On modeling and interpreting the economics of catastrophic climate change’, *The Review of Economics and Statistics* **91**(1), 1–19. 06.
- Weitzman, M. L. (2009b), ‘On modeling and interpreting the economics of catastrophic climate change’, *Review of Economics and Statistics* **91**(1), 1–19.
- Weitzman, M. L. (2010), ‘Ghg targets as insurance against catastrophic climate damages’, *NBER Working Paper* (16136).

Appendix

Part I - Additional Results and Illustrations

A Welfare Results in the Deterministic Setting

Embedded in equation (8) is the social cost of a marginal temperature increase (SCT) in degree Celsius

$$SCT_t(T_{1,t}) = Y_t \xi_0 [(1 - \beta \sigma)^{-1}]_{1,1} \xi_1 \exp(\xi_1 T_{1,t}).$$

The cost of a marginal temperature increase in degree Celsius depends on the prevailing temperature level, unlike the SCC and the transformed temperature state's shadow value. This level-dependence reflects the convexity of damages in temperature. Integrating the shadow value of a temperature increase from pre-industrial to present temperature levels yields the present value welfare cost of the present-day temperature increase

$$\begin{aligned} \Delta W_{USD}^{Temp}{}_{2015}(T_1 \approx 0.77C) &= Y_t \xi_0 [(1 - \beta \sigma)^{-1}]_{1,1} (\exp(\xi_1 T_1) - 1) \\ &\approx \$5 \text{ trillion ,} \end{aligned}$$

or 6% of world output. This value takes into account atmospheric temperature dynamics and the persistence of the global warming. It characterizes the actual cost of having warmed the planet to present temperature levels, which is larger than the annual damage from a given temperature increase, but smaller than the discounted present value of a perpetual temperature increase. The cost does not include that we have already warmed the oceans as well and that the warming is caused by persistent CO₂ emissions that will keep radiative forcing above the pre-industrial level.

The social cost of the present atmospheric CO₂ increase is

$$\Delta W_{USD}^{CO_2}{}_{2015}(M_1 \approx 397ppm) = SCC (M_1 - M_{pre}) \approx \$14 \text{ trillion ,}$$

or 17% of world output. This number reflects the damage already in the pipeline from present atmospheric CO₂. It does not include the CO₂ increase in the oceans or the non-CO₂ greenhouse gases, and the damage is additional to the above cited social cost of the temperature increase that already took place. These numbers illustrate that the welfare cost expected from the present CO₂'s future temperature forcing is significantly higher than the cost of having already heated the planet.

A much discussed geoengineering "solution" to climate change sequesters carbon into the oceans. Engineers are currently exploring mechanisms to extract CO₂ from the exhaustion

pipes of coal power plants, planning to pump it into the deep ocean.²⁶ The gain from such a geoengineering solution is merely the difference between the shadow values of carbon in the different reservoirs. This difference $\varphi_{M,i} - \varphi_{M,1}$ will reappear in the expressions for the welfare loss from carbon flow uncertainty. Appendix E.2 states the closed-form expression for the benefits of pumping a ton of CO₂ into layer i , instead of emitting it into the atmosphere. Appendix C.3 discusses and illustrates the relation between the price of carbon in the different reservoirs. ACE evaluates the welfare gain from pumping a ton of carbon into the upper ocean layer to $57 - 16 = 41$ USD, and to almost the full 57 USD when pumping the carbon into the deep ocean (ACE does not have an explicit damage function for ocean acidification).

B General Capital Depreciation

Equation (3) assumes full capital depreciation. In this appendix, I show how to avoid the full capital depreciation assumption and match observed capital-output ratios through an exogenous adjustment of the capital growth rate. The model extension keeps the structural assumptions that imply a constant investment rate. Under a depreciation rate δ_k the capital accumulation equation (3) changes to

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})] - C_t + (1 - \delta_k)K_t .$$

Defining the consumption rate $x_t = \frac{Y_t[1 - D_t(T_{1,t})]}{C_t}$ and recognizing that $Y_t[1 - D_t(T_{1,t})] - C_t = K_{t+1} - (1 - \delta_k)K_t$ by definition implies

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})](1 - x_t) \left[1 + \frac{1 - \delta_k}{\frac{K_{t+1}}{K_t} - (1 - \delta_k)} \right] .$$

Defining the capital growth rate $g_{k,t} = \frac{K_{t+1}}{K_t} - 1$, I obtain the equivalent equation of motion for capital

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})](1 - x_t) \left[\frac{1 + g_{k,t}}{\delta_k + g_{k,t}} \right] . \quad (\text{B.1})$$

For full depreciation $\delta_k = 1$ the squared bracket is unity and equation (B.1) coincides with equation (3) in the main text. For $\delta_k < 1$ the squared bracket states an empirical correction multiplier larger unity. First, this multiplier can be used to match the model's capital accumulation to the empirical capital accumulation. Second, this multiplier makes the social planner (or representative agent) realize the additional capital value deriving from its persistence beyond its end-of-period value for production. Replacing equation (3) with equation

²⁶Ocean pumping is just one of the strategies considered. Note that CO₂, thanks to its reaction with oxygen, only partially fits back into an oil well. Biochar production corresponds to an injection into soil and the biosphere.

(B.1) does not change the SCC or the formulas for the welfare loss from climate change and uncertainty.

Treating the growth and depreciation correction in squared brackets as exogenous remains an approximation. The extension shows that the model is robust against the immediate criticism of not being able to represent the correct capital evolution and capital output ratio, and against the agent’s neglecting of capital value beyond the time step. The crucial result from and, thus, implicit assumption underlying equations (3) and (B.1) is that the investment rate is independent of the climate states. It is the price to pay for an analytic solution. The remainder of this section shows that this price seems small.

Figure 5 tests ACE’s result (and implicit assumption) that the optimal consumption rate is independent of the climate states. The figure depicts the optimal consumption rate generated by a recursive DICE implementation with an annual time step and, thus, an annual capital decay structure of the usual form (Traeger 2012*b*).²⁷ It also abandons the assumption of logarithmic utility, further stacking the cards against ACE’s assumptions. The first two graphs in the figure depict the control rules for DICE-2013’s $\eta = 1.45$ (inverse of the intertemporal elasticity of substitution). These two graphs state the optimal consumption rate for the years 2025 and 2205. The third graph in the figure depicts the optimal consumption rate for the lower value $\eta = 0.66$ calibrated by the long-run risk literature (see section 4.2).

The qualitative behavior is the same for all graphs in Figure 5. Overall, the figure shows that the optimal consumption rate is largely independent of the climate states (if the vertical axis started at zero the variation of the control rule would be invisible). At current temperature levels, the optimal consumption rate does not depend on the CO₂ concentrations. This result is in accordance with the theoretical result under ACE’s assumption set. However, the optimal consumption rate increases slightly for higher temperatures. It increases by less than a percentage point from no warming to a 3C warming at low CO₂ concentrations. The increase is lower at higher CO₂ concentrations.

The graphs confirm that also in DICE, and in a model with regular annual capital decay structure and not exactly log-utility, the investment rate is not used as a measure of climate change policy. The rate does not respond to the CO₂ concentration, which is a measure of expected warming. Only once the temperature dependent damages set in, the consumption rate slightly increases and the investment rate goes down. Instead of reflecting climate policy, this (minor) climate dependence of the consumption rate reflects a response to the damages incurred: these damages reduce the cake to be split into investment and consumption, then, a slightly higher fraction goes to consumption. This response is lower when CO₂ concentrations

²⁷The recursive implementation based on the Bellman equation solves for the optimal control rule as a function of the states. Thus, solving the model once immediately delivers the full control surface depicted here. This recursive implementation has a slightly simplified climate change model compared to the original DICE model, but matches the Maggic6.0 model, used also as the DICE benchmark, similarly well.

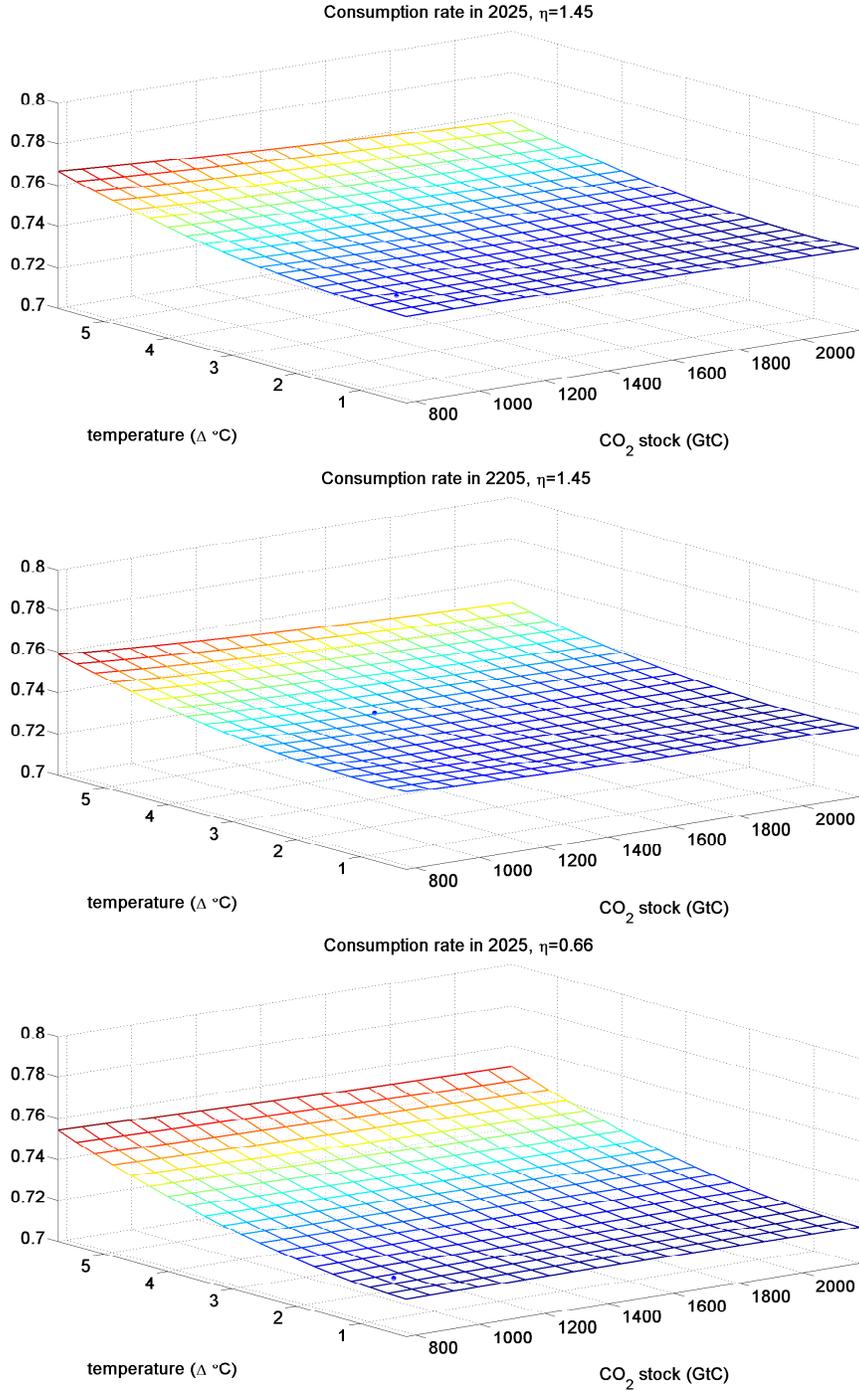


Figure 5: The graphs analyze the climate (in-)dependence of the optimal consumption rate x^* in the wide-spread DICE model, relying on the control rules of the recursive implementation by Traeger (2012b). The first two graphs assume the DICE-2013 value $\eta = 1.45$, the third graph follows the long-run risk literature with $\eta = \frac{2}{3}$. The blue dot in each graph indicates the expected optimal control and prevailing temperature-CO₂ combination along the optimal policy path in the given year.

are high: then the social planner expects high temperatures and damages also in the future and is more hesitant to reduce investment.

C Illustrating the “Climate Matrices”

C.1 A Two Layer Carbon Cycle

In the simple and insightful case of two carbon reservoirs the carbon cycle’s transition matrix is $\Phi = \begin{pmatrix} 1-\delta_{\text{Atm}\rightarrow\text{Ocean}} & \delta_{\text{Ocean}\rightarrow\text{Atm}} \\ \delta_{\text{Atm}\rightarrow\text{Ocean}} & 1-\delta_{\text{Ocean}\rightarrow\text{Atm}} \end{pmatrix}$, where e.g. $\delta_{\text{Atm}\rightarrow\text{Ocean}}$ characterizes the fraction of carbon in the atmosphere transitioning into the ocean in a given time step. The conservation of carbon implies that both columns add to unity: carbon that does not leave a layer (δ_{\rightarrow}) stays ($1 - \delta_{\rightarrow}$). The shadow value becomes

$$\varphi_{M,1} = \beta\varphi_{\tau,1}\sigma^{forc}M_{pre}^{-1}(1 - G\beta)^{-1} \left[1 + G\beta \frac{\delta_{\text{Atm}\rightarrow\text{Ocean}}}{1 - G\beta(1 - \delta_{\text{Ocean}\rightarrow\text{Atm}})} \right]^{-1}.$$

The shadow value becomes less negative if more carbon flows from the atmosphere into the ocean (higher $\delta_{\text{Atm}\rightarrow\text{Ocean}}$). It becomes more negative for a higher persistence of carbon in the ocean (higher $1 - \delta_{\text{Ocean}\rightarrow\text{Atm}}$). These impacts on the SCC are straight forward: the carbon in the ocean is the “good carbon” that does not contribute to the greenhouse effect. In round brackets, I find Proposition ??’s root $(1 - \beta G)^{-1}$ that makes the expression so sensitive to a low rate of pure time preference.

A common approximation of atmospheric carbon dynamics is the equation of motion of the early DICE 1994 model. Here, carbon in excess of preindustrial levels decays as in $M_{1,t+1} = M_{pre} + (1 - \delta_{decay})(M_{1,t} - M_{pre})$. The shadow value formula becomes

$$\varphi_{M,1} = \beta\varphi_{\tau,1}\sigma^{forc}M_{pre}^{-1}(1 - \beta G(1 - \delta_{decay}))^{-1}, \quad (\text{C.1})$$

which misses the long-run carbon impact and the SCC’s sensitivity to pure time preference.

C.2 A Two Layer Atmosphere-Ocean Temperature System

The two layer example of atmosphere-ocean temperature dynamics has the transition matrix $\sigma = \begin{pmatrix} 1-\sigma_1^{up}-\sigma_1^{down} & \sigma_1^{down} \\ \sigma_2^{up} & 1-\sigma_2^{up} \end{pmatrix}$. The corresponding term of the SCC (equation 8) takes the form

$$[(1 - \beta G\sigma)^{-1}]_{11} = \left(1 - \beta G \underbrace{(1 - \sigma_1^{down} - \sigma_1^{up})}_{\text{persistence in atmosphere}} - \frac{\beta^2 \sigma_1^{down} \sigma_1^{up}}{1 - \beta G \underbrace{(1 - \sigma_2^{up})}_{\text{pers. in ocean}}} \right)^{-1}.$$

Persistence of the warming in the atmosphere or in the oceans increases the shadow cost. Persistence of warming in the oceans increases the SCC proportional to the terms σ_1^{down}

routing the warming into the oceans and σ_1^{up} routing the warming back from the oceans into the atmosphere. The discount factor βG accompanies each of these transition coefficients as each of them causes a one period delay. Taking the limit of $\beta G \rightarrow 1$ confirms that (an analogue to) Proposition ?? does not hold for the temperature system

$$\lim_{\beta G \rightarrow 1} \varphi_{\tau,1} = -\xi_0(1 + \varphi_k) [(1 - \boldsymbol{\sigma})^{-1}]_{11} = -\frac{\xi_0(1 + \varphi_k)}{\sigma_1^{up}} \neq \infty. \quad (\text{C.2})$$

As the discount rate approaches zero, the transient temperature dynamics characterized by σ_1^{down} and σ_2^{up} becomes irrelevant for evaluation, and only the weight σ_1^{up} reducing the warming persistence below unity contributes.²⁸

Extending on the “missing time preference sensitivity” in the general case, note that temperature is an intensive variable: it does not scale proportional to mass or volume (as is the case for the extensive variable carbon). The columns of the matrix $\boldsymbol{\sigma}$ do not sum to unity. As a consequence of the mean structure in equation (6), however, the rows in the ocean layers’ transition matrix sum to unity. The first row determining next period’s atmospheric temperature sums to a value smaller than unity: it “misses” the weight that the mean places on radiative forcing. The “missing weight” is a consequence of the permanent energy exchange with outer space. Radiative forcing characterizes a flow equilibrium of incoming and outgoing radiation.

C.3 The Price of Carbon and the Different Reservoirs

The carbon price in the atmosphere is immediately connected to its exchange with the remaining reservoirs. I can also express the shadow value of carbon in any reservoir as the following function of the shadow prices in the remaining reservoirs

$$\varphi_{M,i} = \beta G \frac{\sum_{j \neq i} \varphi_{M,j} \Phi_{j,i} + 1_{i,1} \frac{\varphi_{\tau,1} \sigma_1^{up}}{M_{pre}}}{1 - \beta G \Phi_{i,i}}. \quad (\text{C.3})$$

The carbon price in layer i is the sum of carbon prices in the other layers times the flow coefficient capturing the carbon transition into that other layer (generally only positive for the two adjacent layers). The atmospheric carbon price has as an additional contribution ($1_{i,1}$ denotes the Kronecker delta): the shadow value of the atmospheric temperature increase. Finally, the denominator implies that the carbon price increases over the stated sum according to the persistence $\Phi_{i,i}$ of carbon in that given layer. Equation (C.3) resembles the carbon

²⁸I note that the carbon cycle lacks the reduction in persistence deriving from the forcing weight σ_1^{up} . With this observation, equation (C.2) gives another illustration of the impact of mass conservation in the case of carbon: “ $\sigma_1^{up} \rightarrow 0 \Rightarrow \varphi_{\tau,1} \rightarrow \infty$ ”. Note that in the SCC formula σ_1^{up} cancels, as it simultaneously increases the impact of a carbon change on temperature. This exact cancellation (in the limit $\beta G \rightarrow 1$) is a consequence of the weights σ_1^{up} on forcing and $1 - \sigma_1^{up}$ on atmospheric temperature summing to unity. The result that a warming pulse has a transitional impact and an emission pulse has a permanent impact on the system is independent of the fact that these weights sum to unity.

pricing formula for the carbon decay model discussed in equation (C.1) at the end of section C.1, where the atmospheric carbon persistence is $\Phi_{i,i} = 1 - \delta_{decay}$. The present equation adds the pricing contributions from the other carbon absorbing layers as, unfortunately, the carbon leaving the atmosphere does not decay.

Finally, I illustrate the value of carbon sequestration in equation (E.11) for the case of the two layer carbon cycle introduced in section C.1

$$\Delta W^{seq} = \beta G \varphi_{\tau,1} \sigma_1^{up} M_{pre}^{-1} [1 + G\beta\delta_{\text{Ocean} \rightarrow \text{Atm}} - \beta G(1 - \delta_{\text{Atm} \rightarrow \text{Ocean}})]^{-1}.$$

The value of carbon sequestration into the ocean falls in the stated manner in the transition parameter $\delta_{\text{Ocean} \rightarrow \text{Atm}}$ that captures the carbon diffusion from the ocean back into the atmosphere and increases with the transition parameter $1 - \delta_{\text{Atm} \rightarrow \text{Ocean}}$ that characterizes the persistence of carbon in the atmosphere.

Part II - Solving the Deterministic ACE Model

D The Linear-in-State Model

To obtain the equivalent linear-in-state-system, I first replace capital K_{t+1} by logarithmic capital $k_t \equiv \log K_t$. Second, I replace temperature levels in the atmosphere and the different ocean layers by the transformed exponential temperature states $\tau_{i,t} \equiv \exp(\xi_i T_{i,t})$, $i \in \{1, \dots, L\}$. I collected these transformed temperature states in the vector $\boldsymbol{\tau}_t \in \mathbb{R}^L$. Third, I use the consumption rate $x_t = \frac{C_t}{Y_t[1-D_t(T_t)]}$, rather than absolute consumption, as the consumption-investment control. Only the rate will be separable from the system's states. Finally, I define $a_t = \log A_{0,t}$ and express utility in terms of the consumption rate

$$\begin{aligned} u(C_t(x_t)) &= \log C_t(x_t) = \log x_t + \log Y_t + \log[1 - D_t(T_t)] = \log x_t + a_t \\ &\quad + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} + \nu \log E_t - \xi_0 \exp[\xi_1 T_t] + \xi_0. \end{aligned}$$

The Bellman equation in terms of the transformed state variables is

$$\begin{aligned} V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) &= \max_{x_t, \mathbf{N}_t} \log x_t + a_t + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} \\ &\quad + \nu \log E_t - \xi_0 \tau_t + \xi_0 + \beta V(k_{t+1}, \boldsymbol{\tau}_{t+1}, \mathbf{M}_{t+1}, \mathbf{R}_{t+1}, t+1), \end{aligned} \tag{D.1}$$

and is subject to the following linear equations of motion. The equations of motion for the effective capital stock and the carbon cycle are

$$k_{t+1} = a_t + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} + \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)),$$

$$-\xi_0 \tau_{1,t} + \xi_0 + \log(1 - x_t) \tag{D.2}$$

$$\mathbf{M}_{t+1} = \mathbf{\Phi} \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1 . \tag{D.3}$$

I transform the temperature's equation of motion (6) into a linear system using equation (7)

$$T_{i,t+1} = \frac{1}{\xi_i} \log \left((1 - \sigma_{1,i-1} - \sigma_{1,i+1}) \exp[\xi_i T_{i,t}] \right. \\ \left. + \sigma_{i,i-1} \exp[\xi_i w_i^{-1} T_{i-1,t}] + \sigma_{i,i+1} \exp[\xi_i w_{i+1} T_{i+1,t}] \right) ,$$

the definitions $\sigma_{ii} = 1 - \sigma_{i,i-1} - \sigma_{i,i+1}$, and the requirement $\xi_{t+1} = w_{t+1} \xi_t$. Then the equation is equivalent to

$$\exp(\xi_i T_{i,t+1}) = \sigma_{i,i-1} \exp[\xi_i T_{i,t}] + \sigma_{i,i-1} \exp[\xi_{i-1} T_{i-1,t}] + \sigma_{i,i+1} \exp[\xi_{i+1} T_{i+1,t}] .$$

Using the temperature transformation $\tau_{i,t} = \exp(\xi_i T_{i,t})$ I obtain the linear equations of motion

$$\tau_{i,t+1} = \sigma_{i,i} \tau_{i,t} + \sigma_{i,i-1} \tau_{i-1,t} + \sigma_{i,i+1} \tau_{i+1,t}, i \in \{1, \dots, l\},$$

still using $\sigma_{l,l+1} = 0$ for notational convenience. The first of these equations ($i = 1$) for atmospheric temperature is linear in

$$\tau_{0,t} = \exp[\xi_1 w_1^{-1} T_{0,t}] = \exp \left[\xi_1 \frac{s}{\eta} F_t \right] = \exp \left[\xi_0 \frac{s}{\log 2} \log \frac{M_{1,t} + G_t}{M_{pre}} \right]$$

and has to be linear $M_{1,t}$ to obtain a linear-in-state system (given linearity of the carbon cycle equations). This linearity requires $\xi_0 = \frac{\log 2}{s}$ as stated in the proposition. Then, the individual equations of motion for generalized temperature can be collected into the vector equation

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1 . \tag{D.4}$$

Finally, the equation of motion for the resource stock is

$$\mathbf{R}_{t+1} = \mathbf{R}_t - \mathbf{E}_t^d , \tag{D.5}$$

The underlying constraints are

$$\sum_{i=0}^I N_{i,t} = N_t, N_{i,t} \geq 0,$$

$$\mathbf{R}_t \geq 0 \text{ and } \mathbf{R}_0 \text{ given.}$$

The present paper assumes that the optimal labor allocation has an interior solution and that the scarce resources are stretched over the infinite time horizon along the optimal path, avoiding boundary value complications.

D.0.1 Extended ACE linear in state model

In the extended ACE model of section 2.2 the decision Π_t replaces the consumption rate of the previous section. It determines the share of capital $K_{-1,t} = \Pi_t K_t$ employed in the capital sector. Then utility in a given period is

$$u(C_t) = \log C_t = \log Y_t^{gross} + \log[1 - D_t(T_t)] = \kappa \log \Pi_t + a_t \\ + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} + \nu \log E_t - \xi_0 \exp[\xi_1 T_t] + \xi_0.$$

inducing the Bellman equation

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \max_{x_t, \mathbf{N}_t} \kappa \log \Pi_t + a_t + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} \\ + \nu \log E_t - \xi_0 \tau_t + \xi_0 + \beta V(k_{t+1}, \boldsymbol{\tau}_{t+1}, \mathbf{M}_{t+1}, \mathbf{R}_{t+1}, t+1). \quad (\text{D.6})$$

The only change in the equations of motion affects the effective capital stock

$$k_{t+1} = k_t + \log \left((1 - \delta) + A_{-1,0} (1 - \Pi_t)^{\bar{\kappa}} N_{-1,t}^{1-\bar{\kappa}-\bar{\nu}} \bar{E}_t^{\bar{\nu}} \right) - \bar{\xi}_0 \tau_{1,t} + \bar{\xi}_0,$$

where $\bar{E}_t^{\bar{\nu}} = \frac{A_{-1,t} E_t^{\bar{\nu}}}{A_{-1,0} K^{1-\bar{\kappa}}}$, which is bounded by assumption (detrended).

E Proofs

E.1 Proof of Proposition 1

I start by showing that the affine value function

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t \quad (\text{E.1})$$

solves the above linear-in-state system. The coefficients φ are the shadow value of the respective state variables, and $^\top$ denotes the transpose of a vector of shadow values. The coefficient on the resource stock has to be time-dependent: the shadow value of the exhaustible resource increases (endogenously) over time following the Hotelling rule. The controls in the equations of motion (D.2)-(D.5) are additively separated from the states. Therefore, maximizing the right hand side of the resulting Bellman equation delivers optimal control rules that are independent of the state variables. These controls are functions of the shadow values.

In detail, inserting the value function's trial solution (equation E.1) and the next period states (equations D.2-D.5) into the (deterministic) Bellmann equation (D.6) delivers

$$\begin{aligned}
& \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t = \\
& \max_{x_t, \mathbf{N}_t} \log x_t + \beta \varphi_k \log(1-x_t) + (1 + \beta \varphi_k) \kappa k_t + (1 + \beta \varphi_k) a_t \\
& \quad + (1 + \beta \varphi_k) (1 - \kappa - \nu) \log N_{0,t} \\
& \quad + (1 + \beta \varphi_k) \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) \\
& \quad - (1 + \beta \varphi_k) \xi_0 \tau_{1,t} + (1 + \beta \varphi_k) \xi_0 \\
& \quad + \beta \boldsymbol{\varphi}_M^\top \left(\boldsymbol{\Phi} \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1 \right) \\
& \quad + \beta \boldsymbol{\varphi}_\tau^\top \left(\boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1 \right) \\
& \quad + \beta \boldsymbol{\varphi}_{R,t+1}^\top (\mathbf{R}_t - \mathbf{E}_t^d) \\
& \quad + \beta \varphi_{t+1} \\
& \quad + \lambda_t (N_t - \sum_{i=0}^I N_{i,t})
\end{aligned}$$

Maximizing the right hand side of the Bellman equation over the consumption rate yields

$$\frac{1}{x} - \beta \varphi_k \frac{1}{1-x} = 0 \quad \Rightarrow \quad x^* = \frac{1}{1 + \beta \varphi_k} . \tag{E.2}$$

The labor input into the various sector's depends on the precise assumptions governing the energy sector, i.e., the specification of $g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t))$. For a well-defined energy system, I obtain unique solutions as functions of the technology levels in the energy sector and the shadow values of the endogenous state variables $\mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})$. Knowing these solutions is crucial to determine the precise output path and energy transition under a given policy regime. However, the SCC and, thus, the carbon tax do not depend on these solutions.

Inserting the (general) control rules into the maximized Bellman equation delivers the value function coefficients. In detail, I collect terms that depend on the state variables on

the left hand side of the resulting Bellman equation

$$\begin{aligned}
& \left. \begin{aligned}
& \overbrace{(\boldsymbol{\varphi}_M^\top - \beta \boldsymbol{\varphi}_M^\top \boldsymbol{\Phi} - \beta \varphi_{\tau,1} \frac{\boldsymbol{\sigma}^{forc}}{M_{pre}} \mathbf{e}_1^\top)}^{-\beta b_{det}^M} \mathbf{M}_t + (\boldsymbol{\varphi}_\tau^\top \overbrace{-\beta \boldsymbol{\varphi}_\tau^\top \boldsymbol{\sigma} + (1 + \beta \varphi_k) \xi_0 \mathbf{e}_1^\top}^{-\beta b_{det}^\tau} \boldsymbol{\tau}_t \right\} B \\
& (\varphi_k - (1 + \beta \varphi_k) \kappa) k_t + (\boldsymbol{\varphi}_{R,t}^\top - \beta \boldsymbol{\varphi}_{R,t+1}^\top) \mathbf{R}_t \\
& + \varphi_t = \beta \varphi_{t+1} \\
& \left. \begin{aligned}
& + \log x_t^*(\varphi_k) + \beta \varphi_k \log(1 - x_t^*(\varphi_k)) + (1 + \beta \varphi_k) \xi_0 + (1 + \beta \varphi_k) a_t \\
& + (1 + \beta \varphi_k) (1 - \kappa - \nu) \log \mathbf{N}_{0,t}^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}) \\
& + (1 + \beta \varphi_k) \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}))) \\
& - \beta \boldsymbol{\varphi}_{R,t+1}^\top \mathbf{E}_t^d(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})) \\
& + \beta \varphi_{M,1} \left(\sum_{i=1}^{I^d} E_{i,t}(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})) + E_t^{exo} \right) + \beta \varphi_{\tau,1} \frac{\boldsymbol{\sigma}^{forc}}{M_{pre}} G_t.
\end{aligned} \right\} A \\
& \underbrace{\hspace{15em}}_{a_{det}}
\end{aligned} \tag{E.3}
\end{aligned}$$

The equality holds for all levels of the state variables if and only if the coefficients in front of the state variables vanish, and the evolution of φ_t satisfies the state independent part of the equation. Setting the states' coefficients to zero yields

$$\varphi_k - (1 + \beta \varphi_k) \kappa = 0 \quad \Rightarrow \quad \varphi_k = \frac{\kappa}{1 - \beta \kappa} \tag{E.4}$$

$$\boldsymbol{\varphi}_M^\top - \beta \boldsymbol{\varphi}_M^\top \boldsymbol{\Phi} - \beta \varphi_{\tau,1} \frac{\boldsymbol{\sigma}^{forc}}{M_{pre}} \mathbf{e}_1^\top = 0 \quad \Rightarrow \quad \boldsymbol{\varphi}_M^\top = \frac{\beta \varphi_{\tau,1} \boldsymbol{\sigma}^{forc}}{M_{pre}} \mathbf{e}_1^\top (1 - \beta \boldsymbol{\Phi})^{-1} \tag{E.5}$$

$$\boldsymbol{\varphi}_\tau^\top + (1 + \beta \varphi_k) \xi_0 \mathbf{e}_1^\top - \beta \boldsymbol{\varphi}_\tau^\top \boldsymbol{\sigma} = 0 \quad \Rightarrow \quad \boldsymbol{\varphi}_\tau = -\xi_0 (1 + \beta \varphi_k) \mathbf{e}_1^\top (1 - \beta \boldsymbol{\sigma})^{-1} \tag{E.6}$$

$$\boldsymbol{\varphi}_{R,t}^\top - \beta \boldsymbol{\varphi}_{R,t+1}^\top = 0 \quad \Rightarrow \quad \boldsymbol{\varphi}_{R,t} = \beta^t \boldsymbol{\varphi}_{R,0}. \tag{E.7}$$

The initial values $\boldsymbol{\varphi}_{R,0}^\top$ of the scarce resources depend on the precise evolution of the economy and, thus, depends on assumptions about the energy sector as well as the chosen climate policy. Using the shadow value of log capital in equation (E.2) results in the optimal investment rate $x = 1 - \beta \kappa$. From line (E.8) onwards, the maximized Bellman equation defines recursively the time-dependent affine part of the value function φ_t . Everything discussed in this paper is independent of the process φ_t and only assumes convergence of the value function. For most choices of $g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t))$, the process φ_t has to be solved numerically together with the initial value of shadow price vectors of the scarce resources.

The transformation into the linear-in-state system is unique (up to affine transformations of the states) and the parameter restrictions in Proposition 1 are necessary to obtain linearity. The affine value function solves the system if and only if it is a linear in state system, which completes the proof of Proposition 1.

E.1.1 Extended ACE linear in state model

The linear trial solution for the value function (equation E.1) delivers

$$\begin{aligned}
& \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t = \\
& \max_{x_t, \mathbf{N}_t} \kappa \log \Pi_t + (\kappa + \beta \varphi_k) k_t + a_t + (1 - \kappa - \nu) \log N_{0,t} + \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) - \xi_0 \tau_{1,t} + \xi_0 \\
& \quad \beta \varphi_k [\log ((1 - \delta) + A_{-1,0} (1 - \Pi_t)^{\bar{\kappa}} N_{-1,t}^{1-\bar{\kappa}-\bar{\nu}} \bar{E}_t^{\bar{\nu}}) - \bar{\xi}_0 \tau_{1,t} + \bar{\xi}_0] \\
& \quad + \beta \boldsymbol{\varphi}_M^\top \left(\boldsymbol{\Phi} \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1 \right) \\
& \quad + \beta \boldsymbol{\varphi}_\tau^\top \left(\boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1 \right) \\
& \quad + \beta \boldsymbol{\varphi}_{R,t+1}^\top (\mathbf{R}_t - \mathbf{E}_t^d) \\
& \quad + \beta \varphi_{t+1} \\
& \quad + \lambda_t (N_t - \sum_{i=0}^I N_{i,t}).
\end{aligned}$$

Once again, the controls are separated from the states. Denoting maximized controls by * and sorting by state delivers

$$\begin{aligned}
& \left. \begin{aligned}
& \overbrace{(\boldsymbol{\varphi}_M^\top - \beta \boldsymbol{\varphi}_M^\top \boldsymbol{\Phi} - \beta \varphi_{\tau,1} \frac{\sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top)}^{-\beta b_{det}^M} \mathbf{M}_t + \left(\boldsymbol{\varphi}_\tau^\top \overbrace{-\beta \boldsymbol{\varphi}_\tau^\top \boldsymbol{\sigma}}^{-\beta b_{det}^\tau} + (\xi_0 + \beta \varphi_k \bar{\xi}_0) \mathbf{e}_1^\top \right) \boldsymbol{\tau}_t \\
& (\varphi_k - \kappa - \beta \varphi_k) k_t + (\boldsymbol{\varphi}_{R,t}^\top - \beta \boldsymbol{\varphi}_{R,t+1}^\top) \mathbf{R}_t
\end{aligned} \right\} B_{AK} \\
& + \varphi_t = \beta \varphi_{t+1} \tag{E.8} \\
& \left. \begin{aligned}
& + \log \Pi_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}) + \beta \varphi_k \log \left((1 - \delta) \right. \\
& \left. + A_{-1,0} (1 - \Pi_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}))^{\bar{\kappa}} N_{-1,t}^{*1-\bar{\kappa}-\bar{\nu}} \bar{E}_t^{\bar{\nu}}(\mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})) \right) \\
& + \xi_0 + \beta \varphi_k \bar{\xi}_0 + a_t \\
& + (1 - \kappa - \nu) \log N_{0,t}^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}) \\
& + \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}))) \\
& - \beta \boldsymbol{\varphi}_{R,t+1}^\top \mathbf{E}_t^d(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})) \\
& \underbrace{+ \beta \boldsymbol{\varphi}_{M,1} \left(\sum_{i=1}^{I^d} E_{i,t}(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1})) + E_t^{exo} \right) + \beta \varphi_{\tau,1} \frac{\sigma^{forc}}{M_{pre}} G_t}_{a_{det}}
\end{aligned} \right\} A_{AK}
\end{aligned}$$

The equality holds for all levels of the state variables if and only if the coefficients in front of the state variables vanish. The two resulting shadow values differing from the exogenous

growth model are those for effective capital and temperature

$$\varphi_k(1 - \beta) = \kappa \quad \Rightarrow \quad \varphi_k = \frac{\kappa}{1 - \beta} \quad (\text{E.9})$$

$$\varphi_\tau^\top + (\xi_0 + \beta\varphi_k\bar{\xi}_0)\mathbf{e}_1^\top - \beta\varphi_\tau^\top\boldsymbol{\sigma} = 0 \quad \Rightarrow \quad \varphi_\tau = -(\xi_0 + \beta\varphi_k\bar{\xi}_0)\mathbf{e}_1^\top(1 - \beta\boldsymbol{\sigma})^{-1} \quad (\text{E.10})$$

E.2 Proof of Proposition 2 & Details for Section 3.1

The SCC is the negative of the shadow value of atmospheric carbon expressed in money-measured consumption units. Inserting equation (E.4) for the shadow value of log-capital and (E.6) for the shadow value of atmospheric temperature (first entry of the vector) into equation (E.5) characterizing the shadow value of carbon in the different reservoirs delivers

$$\varphi_M^\top = -\xi_0 \left(1 + \beta \frac{\kappa}{1 - \beta\kappa}\right) [(1 - \beta\boldsymbol{\sigma})^{-1}]_{1,1} \frac{\beta\sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top (\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}.$$

As a consequence of logarithmic utility, this marginal welfare change translates into a consumption change as $du = \frac{1}{c}dc = \frac{1}{xY}dc \Rightarrow dc = (1 - \beta\kappa)Ydu$. Thus, observing that $(1 + \beta \frac{\kappa}{1 - \beta\kappa}) = \frac{1}{1 - \beta\kappa}$, the SCC is

$$SCC = -(1 - \beta\kappa)Y_t\varphi_{M,1} = Y_t \xi_0 [(1 - \beta\boldsymbol{\sigma})^{-1}]_{1,1} \frac{\beta\sigma^{forc}}{M_{pre}} [(\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}]_{1,1}.$$

The social cost of an atmospheric temperature increase follows similarly from the shadow value of the generalized temperature state $\tau_{1,t}$

$$SC\tau = -(1 - \beta\kappa)Y_t\varphi_{\tau,1} = Y_t \xi_0 [(1 - \beta\boldsymbol{\sigma})^{-1}]_{1,1}.$$

A marginal increase in generalized temperature relates to a temperature increase in degree Celsius as $d\tau_{1,t} = \xi_1 \exp(\xi_1 T_{1,t}) dT_{1,t}$ implying the social cost of a temperature increase in degree Celsius of

$$SCT(T_{1,t}) = Y_t \xi_0 [(1 - \beta\boldsymbol{\sigma})^{-1}]_{1,1} \xi_1 \exp(\xi_1 T_{1,t}).$$

The welfare loss from present temperature increase integrates this formula from 0 to the present temperature.

Pumping a ton of CO₂ into layer i , instead of emitting it into the atmosphere, results in the welfare gain

$$\Delta W^{seq} = \varphi_{M,i} - \varphi_{M,1} = \frac{\beta G \varphi_{\tau,1} \sigma^{forc}}{M_{pre}} \left([(\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}]_{1,i} - [(\mathbf{1} - \beta\boldsymbol{\Phi})^{-1}]_{1,1} \right). \quad (\text{E.11})$$

The bracket on the right hand side captures the discounted sum of the differences in the amount of carbon prevailing in the atmosphere over time when an emission unit is injected

into layer i instead of the atmosphere. This intuition is more easily observed using the Neumann series for the expression:

$$\Delta W^{seq} = \frac{\beta G \varphi_{\tau,1} \sigma_1^{up}}{M_{pre}} \left(\beta G [\Phi_{1,i} - \Phi_{1,1}] + \sum_{n=2}^{\infty} \sum_{j,l} (\beta G)^n \Phi_{1,j} (\Phi^{n-2})_{j,l} [\Phi_{l,i} - \Phi_{l,1}] \right).$$

The first term in the brackets captures the difference between carbon flow from the ocean into the atmosphere $\Phi_{1,i}$ and the persistence of carbon in the atmosphere $\Phi_{1,1}$. The second term captures the fraction of carbon reaching the atmosphere after n periods if the carbon initially enters ocean layer i as opposed to entering the atmosphere directly (read right to left). The matrix entry $(\Phi^{n-2})_{j,l}$ captures the overall carbon flow and persistence from layer l to j after $n - 2$ periods. It approaches the stationary distribution given by its (right) eigenvectors (in all columns). In the DICE carbon cycle, the value of sequestering carbon into the intermediate ocean and biosphere corresponding is \$41 per ton and the value of pumping carbon into the deep ocean is \$56 per ton.²⁹ Appendix C.3 illustrates equation (E.11) for a two layer carbon cycle and discusses more generally the relation between carbon prices in different reservoirs.

E.3 Proof of Proposition 3

Part (i). Mass conservation of carbon implies that the columns of Φ add to unity. In consequence, the vector with unit entry in all dimensions is a left and, thus, right eigenvector. The corresponding eigenvalue is one and the determinant of $\mathbf{1} - \beta G \Phi$ has the root $1 - \beta G$. It follows from Cramer's rule (or as an application of the Cayley-Hamilton theorem) that the entries of the matrix $(\mathbf{1} - \beta G \Phi)^{-1}$ are proportional to $(1 - \beta G)^{-1}$.

Part (ii), SCC of extended DICE. Inserting the equations (E.9) for the shadow value of log-capital and (E.10) for the shadow value of atmospheric temperature (first entry of the vector) into the earlier equation (E.5) delivers

$$\begin{aligned} \varphi_M^\top &= -(\xi_0 + \beta \varphi_k \bar{\xi}_0) [(1 - \beta \sigma)^{-1}]_{1,1} \frac{\beta \sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top (\mathbf{1} - \beta \Phi)^{-1} \\ &= -\left(\xi_0 + \bar{\xi}_0 \frac{\beta \kappa}{1 - \beta} \right) [(1 - \beta \sigma)^{-1}]_{1,1} \frac{\beta \sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top (\mathbf{1} - \beta \Phi)^{-1}. \end{aligned}$$

As a consequence of logarithmic utility, this marginal welfare change translates into a consumption change as $du = \frac{1}{\epsilon} dc$ and the SCC is

$$SCC = -C_t \varphi_{M,1} = C_t \left(\xi_0 + \bar{\xi}_0 \frac{\beta \kappa}{1 - \beta} \right) [(1 - \beta \sigma)^{-1}]_{1,1} \frac{\beta \sigma^{forc}}{M_{pre}} [(\mathbf{1} - \beta \Phi)^{-1}]_{1,1}.$$

²⁹Note that the present model does not explicitly model damages from ocean acidification, which would be an interesting and feasible extension.

Alternatively, the shadow value of atmospheric carbon can be transformed into capital units using the shadow value of capital (rather than log capital), which is $\varphi_K = \frac{\varphi_k}{K}$ (as $dk = d \log K = \frac{dK}{K}$). Then

$$\begin{aligned}
SCC &= -K_t \frac{\varphi_{M,1}}{\varphi_k} = K_t \left(\frac{\xi_0}{\varphi_k} + \bar{\xi}_0 \beta \right) [(1 - \beta \boldsymbol{\sigma})^{-1}]_{1,1} \frac{\beta \boldsymbol{\sigma}^{forc}}{M_{pre}} [(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1} \\
&= K_t \left(\frac{\xi_0}{\varphi_k} + \bar{\xi}_0 \beta \right) [(1 - \beta \boldsymbol{\sigma})^{-1}]_{1,1} \frac{\beta \boldsymbol{\sigma}^{forc}}{M_{pre}} [(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1} \\
&= \frac{K_t}{\kappa} ((1 - \beta) \xi_0 + \beta \kappa \bar{\xi}_0) [(1 - \beta \boldsymbol{\sigma})^{-1}]_{1,1} \frac{\beta \boldsymbol{\sigma}^{forc}}{M_{pre}} [(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1} .
\end{aligned}$$

SCC of base ACE transforms into capital equivalents accordingly.

Part(iii). To come.

Part II - Uncertainty

F Appendix to Section 4: General Results

F.1 Equivalence to Epstein-Zin-Weil Utility and Illustration of Risk Aversion

This section presents a quantitative illustration of the adopted risk aversion and derives the equivalence to Epstein-Zin-Weil preferences. I start by showing the equivalence of the Bellman equation (9) to the wide-spread formulation of recursive utility going back to Epstein & Zin (1991) and Weil (1990). Keeping isoelastic risk aggregation and using the logarithmic special case for intertemporal aggregation reflecting GAUVAL's intertemporal elasticity of unity, the usual formulation reads

$$V_t^* = \exp \left((1 - \beta) \log c_t + \beta \log \left[\mathbb{E}_t V_{t+1}^* \right]^{\frac{1}{\alpha^*}} \right). \quad (\text{F.1})$$

Defining $V_t = \frac{\log V_t^*}{1-\beta}$ and rearranging equation (F.1) delivers

$$V_t = \log c_t + \frac{\beta}{1-\beta} \log \left[\mathbb{E}_t \exp \left((1 - \beta) V_{t+1} \right)^{\frac{1}{\alpha^*}} \right].$$

Defining $\alpha = (1 - \beta)\alpha^*$ and pulling the risk aversion coefficient α^* of the Epstein-Zin setting to the front of the logarithm and into the exponential yields equation (9) stated in the text.

Figure 6 illustrates the quantitative implications of a choice of risk aversion $\text{RRA} = 1 - \alpha$ in the model.³⁰ In the baseline, an agent consumes a constant level \bar{c} in perpetuity. In a coin toss lottery, she loses 5% of her consumption in the upcoming decade (left) or 25% (right) in case of tails (probability $1/2$). The graph presents, as a function of her risk aversion RRA, the percentage gain over the baseline that the agent requests if heads comes up to be indifferent between the lottery and the baseline. It is important to realize that these losses and gains are direct consumption changes. The numeric illustrations in the paper are based on the range $\text{RRA}^* = 1 - \alpha^* \in [6, 9.5]$ found in the long-run risk literature. The bounds translate approximately into $\alpha = (1 - \beta)\alpha^* \in \{1, 1.5\}$ in the present model's equation (9) and into $\text{RRA} \in \{2, 2.5\}$ in Figure 6.

³⁰I directly illustrate risk aversion for the choice of $1 - \alpha$ as opposed to Epstein-Zin's $1 - \alpha^* = 1 - \frac{\alpha}{1-\beta}$. This illustration is independent of time preference. A similar time preference independent illustration of Epstein-Zin's $1 - \alpha^*$ would involve a lottery over infinite consumption streams. The argument why $1 - \alpha^*$ as opposed to $1 - \alpha$ would be time preference invariant relies on the idea that the lottery payoffs in the current period have less significance for a more patient agent.

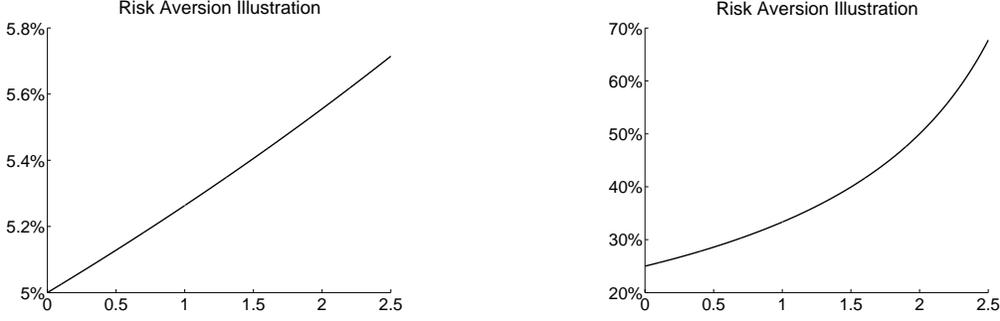


Figure 6: The graphs illustrate the relation between the risk aversion $RRA = 1 - \alpha$ and the relative consumption gains and losses that leave an agent indifferent to her original position. With probability $1/2$, the agent loses 5% of her decadal consumption (left) or 25% (right). The graphs show how much of a relative consumption gain she requires for being indifferent to her initial deterministic position under different degrees of risk aversion.

F.2 Proof of Proposition 4

Inserting an affine trial solution of the value function into the Bellman equation (9) and using the same transformations as in the deterministic case delivers

$$\begin{aligned}
 & \varphi_k k_t + \varphi_M^\top \mathbf{M}_t + \varphi_\tau^\top \boldsymbol{\tau}_t + \varphi_{R,t}^\top \mathbf{R}_t + \varphi_t + \varphi_I^\top \mathbf{I}_t = \\
 & \max_{x_t, \mathbf{N}_t} \log x_t + \beta \varphi_k \log(1 - x_t) + (1 + \beta \varphi_k) \kappa k_t + (1 + \beta \varphi_k) a_t \\
 & \quad + (1 + \beta \varphi_k) (1 - \kappa - \nu) \log N_{0,t} \\
 & \quad + (1 + \beta \varphi_k) \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) + \varphi_{t+1} \\
 & \quad - (1 + \beta \varphi_k) \xi_0 \tau_{1,t} + (1 + \beta \varphi_k) \xi_0 + \varphi_{R,t+1}^\top \mathbf{R}_{t+1} \\
 & \quad + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\varphi_M^\top \mathbf{M}_{t+1} + \varphi_\tau^\top \boldsymbol{\tau}_{t+1} + \varphi_I^\top \mathbf{I}_{t+1} \right) \right] \right).
 \end{aligned}$$

Using assumption (11) on the conditional expectation $\mathbf{X}_t = (\mathbf{M}_t, \boldsymbol{\tau}_t, \mathbf{I}_t)$ with $\mathbf{z} = (\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top)$ yields

$$\begin{aligned}
& \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t + \boldsymbol{\varphi}_I^\top \mathbf{I}_t = \\
& \max_{x_t, \mathbf{N}_t} \log x_t + \beta\varphi_k \log(1-x_t) + (1 + \beta\varphi_k)\kappa k_t + (1 + \beta\varphi_k)a_t \\
& \quad + (1 + \beta\varphi_k)(1 - \kappa - \nu) \log N_{0,t} \\
& \quad + (1 + \beta\varphi_k)\nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) \\
& \quad - (1 + \beta\varphi_k)\xi_0 \tau_{1,t} + (1 + \beta\varphi_k)\xi_0 \\
& \quad + \beta\boldsymbol{\varphi}_{R,t+1}^\top (\mathbf{R}_t - \mathbf{E}_{1,t}) + \beta\varphi_{t+1} + \lambda_t (N_t - \sum_{i=0}^I N_{i,t}) \\
& \quad + \frac{\beta}{\alpha} \left(a(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top) + \sum_{i=1}^N b_i(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top) X_{t,i} \right).
\end{aligned}$$

Maximizing the right hand side of the Bellman equation implies the same optimal consumption rate $x^* = \frac{1}{1+\beta\varphi_k}$ as in the deterministic case, and a set of general control rules for the energy sector inputs $\mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}, \boldsymbol{\varphi}_I)$. Inserting the optimal control rules and collecting the state-dependent terms (ordered by state) on the left hand side of the equality yields

$$\begin{aligned}
& (\varphi_k - (1 + \beta\varphi_k)\kappa)k_t + \sum_{i=1}^m \left(\varphi_{M,i} - \frac{\beta}{\alpha} b_i^M(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top) \right) M_{i,t} \\
& \quad + \sum_{i=1}^l \left(\varphi_{\tau,i} + (1 + \beta\varphi_k)\xi_0 \delta_{i,1} - \frac{\beta}{\alpha} b_i^\tau(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top) \right) \tau_{i,t} \\
& \quad + (\boldsymbol{\varphi}_{R,t}^\top - \beta\boldsymbol{\varphi}_{R,t+1}^\top) \mathbf{R}_t + \varphi_t \\
& \quad + \sum_{i=1}^{N-l-m} \left(\varphi_{I,i} - \frac{\beta}{\alpha} b_i^I(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top) \right) I_{i,t} \\
& = \log x_t^*(\varphi_k) + \beta\varphi_k \log(1-x_t^*(\varphi_k)) + (1 + \beta\varphi_k)a_t + (1 + \beta\varphi_k)\xi_0 \\
& \quad + (1 + \beta\varphi_k)(1 - \kappa - \nu) \log \mathbf{N}_{0,t}^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}, \boldsymbol{\varphi}_I) \\
& \quad + (1 + \beta\varphi_k)\nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}, \boldsymbol{\varphi}_I))) \\
& \quad - \beta\boldsymbol{\varphi}_{R,t+1}^\top \mathbf{E}_t^d(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \boldsymbol{\varphi}_M, \boldsymbol{\varphi}_{R,t+1}, \boldsymbol{\varphi}_I)) \\
& \quad + \beta\varphi_{t+1} + \frac{\beta}{\alpha} a(\alpha\boldsymbol{\varphi}_M^\top, \alpha\boldsymbol{\varphi}_\tau^\top, \alpha\boldsymbol{\varphi}_I^\top),
\end{aligned}$$

where $(b_1^M, \dots, b_m^M, b_1^\tau, \dots, b_l^\tau, b_1^I, \dots, b_{N-l-m}^I) = (b_1, \dots, b_N)$ and $\delta_{i,j}$ denotes the Kronecker-delta (one if $i = j$ and zero otherwise). The trial solution solves the stochastic optimization problem if (and only if) all the coefficients in front of the state variables vanish. The coefficient on log capital results in $\varphi_k = \frac{\kappa}{1-\beta\kappa}$ (see as well equation E.4). The coefficient in front of the resource vector implies Hotelling's rule $\boldsymbol{\varphi}_{R,t} = \beta^t \boldsymbol{\varphi}_{R,0}$ (see as well equation E.7). The time dependent affine shadow value φ_t can be chosen to set the right hand side of the Bellman

equation to zero (thereby measuring the state-independent welfare contribution). Thus, the trial solution solves the dynamic programming problem if and only if the shadow values solve the equations stated in the proposition (eliminating the coefficients of the remaining states).

G Appendix to Section 5: Results under Uncertainty

G.1 Proof of Proposition 5

The autoregressive gamma process by Gouriéroux & Jasiak (2006) is as a Poisson mixture of gamma distributions. Denoting by X_t the state of an autoregressive gamma process the one-step-ahead distribution is

$$\frac{X_{t+1}}{c} | (Z, X_t) \sim \text{gamma}(\nu_t + Z), \text{ where } Z | X_t \sim \text{Poisson} \left(\frac{\gamma X_t}{c} \right).$$

for $c, \gamma, \nu_t > 0$ in all periods. The random variable Z is drawn from a Poisson distribution and modulates the shape parameter of the standard gamma distribution (with scale c). The expectation and variance of this process are

$$\mathbb{E}(X_{i,t+1} | X_t) = \nu_t c + \gamma X_t$$

$$\text{Var}(X_{i,t+1} | X_t) = \nu_t c^2 + 2c\gamma X_t.$$

and the cumulant generating function is

$$G_{X_{t+1}}(u) = \log [\mathbb{E}(\exp(uX_{t+1}) | X_t)] = -\nu_t \log(1 - uc) + \frac{u}{1-uc} \gamma X_t.$$

Applying the model to the temperature-carbon feedback, I chose

$$\nu_t = \frac{\delta_\tau}{c} \left(\frac{M_{1,t+G_t}}{M_{pre}} - \eta_\tau \right),$$

which results in the expectation and variance cited as equations (16) and (17) in the main text. To apply Proposition 4, I calculate (one over α times) the cumulant generation function

of $\mathbf{X}_t = (\mathbf{M}_t, \boldsymbol{\tau}_t, \mathbf{I}_t)$ with $\mathbf{I}_t = (x_t^M, \sigma_t^M, x_t^\tau, z_t)$ and $\mathbf{z} = \alpha \boldsymbol{\varphi}^\top$

$$\begin{aligned} \frac{1}{\alpha} \log (\mathbb{E} \exp (\alpha \boldsymbol{\varphi}^\top \mathbf{X}_{t+1}) | \mathbf{X}_t) &= \boldsymbol{\varphi}_M^\top \boldsymbol{\Phi} \mathbf{M}_t + \dots + (\varphi_{M,1} - \varphi_{M,2}) x_t^M + \frac{\alpha}{2} (\varphi_{M1} - \varphi_{M2})^2 \sigma_t^{M2} \quad (\text{G.1}) \\ &+ \boldsymbol{\varphi}_\tau^\top \boldsymbol{\sigma} \boldsymbol{\tau}_t + \frac{\sigma^{forc}}{M_{pre}} \varphi_{\tau,1} M_{1,t} - h \varphi_{\tau,1} \gamma_\tau^x z_t - h \varphi_{\tau,1} (\delta_\tau - \epsilon(c)) \frac{M_{1,t}}{M_{pre}} + \dots \\ &+ \varphi_x^M \gamma_M^x x_t^M + \frac{\alpha}{2} \varphi_x^{M2} \delta_M^{Mx2} \frac{M_{1,t}}{M_{pre}} + \frac{\alpha}{2} \varphi_x^{M2} \delta_M^{\sigma x2} \sigma_t^{M2} \\ &+ \varphi_\sigma^M \gamma_M^\sigma \sigma_t^{M2} + \dots + \varphi_\sigma^M \delta_M^{M\sigma} \frac{M_{1,t}}{M_{pre}} \\ &- \frac{\delta_\tau}{\alpha c} \frac{M_{1,t}}{M_{pre}} \log (1 - \alpha [\varphi_x^\tau + h \varphi_{\tau,1}] c) + \frac{\varphi_x^\tau + h \varphi_{\tau,1}}{1 - \alpha [\varphi_x^\tau + h \varphi_{\tau,1}] c} \gamma_\tau^x x_t^\tau \\ &+ \varphi_z \gamma_\tau^x z_t + \varphi_z (\delta_\tau - \epsilon(c)) \frac{M_{1,t}}{M_{pre}} + \dots \end{aligned}$$

I abbreviate by “...” affine terms that are independent of the states. Sorting the r.h.s. of equation (G.1) by states identifies Proposition 4’s linear terms $b_i^M(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top)$, $b_i^\tau(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top)$, and $b_i^I(\alpha \boldsymbol{\varphi}_M^\top, \alpha \boldsymbol{\varphi}_\tau^\top, \alpha \boldsymbol{\varphi}_I^\top)$ where $\boldsymbol{\varphi}_I = (\varphi_x^M, \varphi_x^\tau, \varphi_\sigma^M, \varphi_\sigma^\tau)$. Then, Proposition 4 implies the following system of equations for the shadow values

$$\begin{aligned} \boldsymbol{\varphi}_M^\top &= \beta \boldsymbol{\varphi}_M^\top \boldsymbol{\Phi} + \beta \left(\frac{\sigma^{forc}}{M_{pre}} \varphi_{\tau,1} + \frac{\alpha}{2} \frac{\delta_M^{Mx2}}{M_{pre}} \varphi_x^{M2} + \frac{\delta_M^{M\sigma}}{M_{pre}} \varphi_\sigma^M + \frac{1}{M_{pre}} (\varphi_z^\tau - h \varphi_{\tau,1}) (\delta_\tau - \epsilon(c)) \right. \\ &\quad \left. - \frac{\delta_\tau}{M_{pre}} \frac{\log(1 - \alpha c (\varphi_x^\tau + h \varphi_{\tau,1}))}{\alpha c} \right) \mathbf{e}_1^\top \quad (\text{G.2}) \end{aligned}$$

$$\boldsymbol{\varphi}_\tau^\top = \beta \boldsymbol{\varphi}_\tau^\top \boldsymbol{\sigma} - (1 + \beta \varphi_k) \xi_0 \mathbf{e}_1^\top \quad (\text{G.3})$$

$$\varphi_z^\tau = \beta (\varphi_z^\tau - h \varphi_{\tau,1}) \gamma_\tau^x \quad (\text{G.4})$$

$$\varphi_x^\tau = \beta \frac{\varphi_x^\tau + h \varphi_{\tau,1}}{1 - \alpha c (\varphi_x^\tau + h \varphi_{\tau,1})} \gamma_\tau^x \quad (\text{G.5})$$

$$\varphi_x^M = \beta (\varphi_{M,1} - \varphi_{M,2}) + \beta \varphi_x^M \gamma_M^x \quad (\text{G.6})$$

$$\varphi_\sigma^M = \beta \frac{\alpha}{2} \left((\varphi_{M,1} - \varphi_{M,2})^2 + \delta_M^{\sigma x2} \varphi_x^{M2} \right) + \beta \varphi_\sigma^M \gamma_M^\sigma \quad (\text{G.7})$$

Temperature related shadow values:

The temperature’s shadow value is as before by equation (G.3)

$$\boldsymbol{\varphi}_\tau^\top = -(1 + \beta \varphi_k) \xi_0 \mathbf{e}_1^\top (\mathbf{I} - \beta \boldsymbol{\sigma})^{-1}.$$

The feedback operates through the carbon’s shadow value and through the persistent shock shadow value φ_x^τ for which equation (G.5) delivers the quadratic equation

$$\begin{aligned} \varphi_x^\tau - \alpha c \varphi_x^{\tau2} - \alpha c h \varphi_{\tau,1} \varphi_x^\tau &= \beta \varphi_x^\tau \gamma_\tau^x + \beta h \varphi_{\tau,1} \gamma_\tau^x \\ \Leftrightarrow \underbrace{\alpha c}_{\equiv \tilde{a}} \varphi_x^{\tau2} + \underbrace{(\beta \gamma_\tau^x + \alpha c h \varphi_{\tau,1} - 1)}_{\equiv \tilde{b}} \varphi_x^\tau + \underbrace{\beta h \varphi_{\tau,1} \gamma_\tau^x}_{\equiv \tilde{c}} &= 0 \end{aligned}$$

Instead of using the common abc -formula I use the solution arrived at by Mullers method, which solves $\tilde{a}x^2 + \tilde{b}x + \tilde{c} = 0$ by the roots $x = \frac{-2\tilde{c}}{\tilde{b} \pm \sqrt{\tilde{b}^2 - 4\tilde{a}\tilde{c}}}$. The solution is advantageous because it yields a valid root for the case $\tilde{a} = 0$, which corresponds to the deterministic case.³¹ Then

$$\begin{aligned}\varphi_x^\tau &= \frac{-2\tilde{c}}{\tilde{b} \pm \sqrt{\tilde{b}^2 - 4\tilde{a}\tilde{c}}} = \varphi_{\tau,1} \frac{2\beta h\gamma_\tau^x}{(1 - \beta\gamma_\tau^x - \alpha ch\varphi_{\tau,1}) \pm \sqrt{(1 - \beta\gamma_\tau^x - \alpha ch\varphi_{\tau,1})^2 - 4\alpha ch\varphi_{\tau,1}\beta\gamma_\tau^x}} \\ &= \frac{\beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x} \frac{2}{\underbrace{1 - \frac{\alpha ch\varphi_{\tau,1}}{1 - \beta\gamma_\tau^x} \pm \sqrt{\left(1 - \frac{\alpha ch\varphi_{\tau,1}}{1 - \beta\gamma_\tau^x}\right)^2 - 4\frac{\alpha ch\varphi_{\tau,1}}{1 - \beta\gamma_\tau^x} \frac{\beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x}}}_{\equiv T}} h\varphi_{\tau,1}\end{aligned}\quad (\text{G.8})$$

To identify the economically meaningful root, I take $c \rightarrow 0$. The negative root diverges and identifies the positive root as the correct root (the root with $+\sqrt{\quad}$). The correct deterministic limit delivers $\varphi_x^\tau \rightarrow \varphi_{\tau,1} \frac{\beta h\gamma_\tau^x}{(1 - \beta\gamma_\tau^x)}$ for $c \rightarrow 0$. The shadow value in the deterministic limit coincides with the (negative of the) shadow value φ_z^τ that results from equation (G.5) as

$$\varphi_z^\tau = -\frac{\beta h\gamma_\tau^x}{1 - \beta\gamma_\tau^x} \varphi_{\tau,1}.$$

Carbon-flow uncertainty:

Equation (G.2) delivers the shadow value vector equation

$$\begin{aligned}\varphi_M^\top &= \beta \left(\frac{\sigma^{forc}}{M_{pre}} \varphi_{\tau,1} + \frac{\alpha \delta_M^{Mx^2}}{2 M_{pre}} \varphi_x^{M^2} + \frac{\delta_M^{M\sigma}}{M_{pre}} \varphi_\sigma^M + \frac{1}{M_{pre}} (\varphi_z^\tau - h\varphi_{\tau,1}) (\delta_\tau - \epsilon(c)) \right. \\ &\quad \left. - \frac{\delta_\tau}{M_{pre}} \frac{\log(1 - \alpha c (\varphi_x^\tau + h\varphi_{\tau,1}))}{\alpha c} \right) [(\mathbb{I} - \beta\Phi)^{-1}]_{1, \cdot}.\end{aligned}\quad (\text{G.9})$$

Dividing the second through the first shadow value entry I obtain

$$\varphi_{M,2} = \underbrace{\frac{[(\mathbb{I} - \beta\Phi)^{-1}]_{1,2}}{[(\mathbb{I} - \beta\Phi)^{-1}]_{1,1}}}_{\equiv r} \varphi_{M,1}.\quad (\text{G.10})$$

Equation (G.6) delivers the shadow value

$$\varphi_x^M = \frac{\beta}{1 - \gamma_M^x \beta} (\varphi_{M,1} - \varphi_{M,2}) = \underbrace{\frac{\beta}{1 - \gamma_M^x \beta}}_{\equiv A} (1 - r) \varphi_{M,1},\quad (\text{G.11})$$

where the second equality uses equation (G.10). Substituting these results into equation

³¹The common abc -formula yields a fraction $\frac{0}{0}$ for $\tilde{a} = 0$. Having a well-defined root for the deterministic special case permits me to connect the uncertain SCC directly to the deterministic SCC.

(G.7) delivers

$$\varphi_\sigma^M = \beta \frac{\alpha (\varphi_{M,1} - \varphi_{M,2})^2 + \delta_M^{\sigma x^2} \varphi_x^{M^2}}{1 - \gamma_M^\sigma \beta} = \beta \underbrace{\frac{\alpha (1-r)^2 + \delta_M^{\sigma x^2} A^2}{1 - \gamma_M^\sigma \beta}}_{\equiv B} \varphi_{M,1}^2. \quad (\text{G.12})$$

Inserting equation (G.11) and (G.12) into the atmospheric shadow value component of equation (G.9) results in the quadratic equation

$$\begin{aligned} \varphi_{M,1} &= \beta \left(\frac{\alpha \delta_M^{Mx^2}}{2 M_{pre}} \varphi_x^{M^2} + \frac{\delta_M^{M\sigma}}{M_{pre}} \varphi_\sigma^M \right) [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \\ &+ \frac{\beta}{M_{pre}} \left(\sigma^{forc} \varphi_{\tau,1} + (\varphi_z^\tau - h \varphi_{\tau,1})(\delta_\tau - \epsilon(c)) - \delta_\tau \frac{\log(1 - \alpha c(\varphi_x^\tau + h \varphi_{\tau,1}))}{\alpha c} \right) [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \\ &= \beta \underbrace{\left(\frac{\alpha \delta_M^{Mx^2}}{2 M_{pre}} A^2 + \frac{\delta_M^{M\sigma}}{M_{pre}} B \right)}_{\equiv \hat{a}} [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \varphi_{M,1}^2 \\ &+ \underbrace{\frac{\beta}{M_{pre}} \left(\sigma^{forc} \varphi_{\tau,1} + (\varphi_z^\tau - h \varphi_{\tau,1})(\delta_\tau - \epsilon(c)) - \delta_\tau \frac{\log(1 - \alpha c(\varphi_x^\tau + h \varphi_{\tau,1}))}{\alpha c} \right)}_{\equiv \hat{c}} [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1}. \end{aligned}$$

Using once more the quadratic formula deriving from Muller's method I obtain the solution

$$\varphi_{M,1} = \frac{2\hat{c}}{1 \pm \sqrt{1 - 4\hat{a}\hat{c}}}$$

and once again the positive root is the one that is economically meaningful as it converges for $\hat{a} = 0$ to the correct solution (including the deterministic special case if all uncertainty is absent). I transform the expression for $\varphi_{M,1}$ and, in the last step, do a first order Taylor approximation in both numerator and denominator

$$\varphi_{M,1} = \frac{2\hat{c}}{1 + \sqrt{1 - 4\hat{a}\hat{c}}} = \hat{c} \left(1 + \underbrace{\frac{1 - \sqrt{1 - 4\hat{a}\hat{c}}}{1 + \sqrt{1 - 4\hat{a}\hat{c}}}}_{\equiv \theta_M} \right) \approx \hat{c} \left(1 + \frac{\hat{a}\hat{c}}{1 - \hat{a}\hat{c}} \right), \quad (\text{G.13})$$

where the approximation is first order around $\hat{a}\hat{c} = 0$ in both numerator and denominator.

The term \hat{a} is

$$\begin{aligned}
\hat{a} &= \beta \left(\frac{\alpha \delta_M^{Mx^2}}{2 M_{pre}} A^2 + \frac{\delta_M^{M\sigma}}{M_{pre}} B \right) [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \\
&= \beta \frac{\alpha}{2} \frac{1}{M_{pre}} \left[\left(\frac{\beta \delta_M^{Mx}}{1 - \gamma_M^x \beta} \right)^2 (1-r)^2 \right. \\
&\quad \left. + \frac{\beta \delta_M^{M\sigma}}{1 - \gamma_M^\sigma \beta} \left((1-r)^2 + \left(\frac{\beta \delta_M^{\sigma x}}{1 - \gamma_M^x \beta} \right)^2 (1-r)^2 \right) \right] [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \\
&= \frac{\alpha}{2} \frac{\beta}{M_{pre}} \left[A_M^{M \rightarrow x^2} + A_M^{M \rightarrow \sigma} A_M^{\sigma \rightarrow x^2} + A_M^{M \rightarrow \sigma} \right] (1-r)^2 [(\mathbb{I} - \beta \Phi)^{-1}]_{1,1} \\
&= \frac{\alpha}{2} \frac{\beta}{M_{pre}} \left[A_M^{M \rightarrow x^2} + A_M^{M \rightarrow \sigma} A_M^{\sigma \rightarrow x^2} + A_M^{M \rightarrow \sigma} \right] \frac{([\mathbb{I} - \beta \Phi]^{-1})_{1,1} - ([\mathbb{I} - \beta \Phi]^{-1})_{1,2}}{([\mathbb{I} - \beta \Phi]^{-1})_{1,1}} \tag{G.14}
\end{aligned}$$

$$\text{with } A_M^{M \rightarrow x} = \frac{\delta_M^{Mx} \beta}{1 - \gamma_M^x \beta}, \quad A_M^{M \rightarrow \sigma} = \frac{\delta_M^{M\sigma} \beta}{1 - \gamma_M^\sigma \beta}, \quad A_M^{\sigma \rightarrow x} = \frac{\delta_M^{\sigma x} \beta}{1 - \gamma_M^x \beta}. \tag{G.15}$$

Temperature-carbon feedback:

Evaluating the term \hat{c} requires the evaluation of

$$\varphi_z^\tau - h \varphi_{\tau,1} = - \left(\frac{\beta \gamma_\tau^x}{1 - \beta \gamma_\tau^x} + 1 \right) h \varphi_{\tau,1} = - \frac{h}{1 - \beta \gamma_\tau^x} \varphi_{\tau,1} = - \bar{h} \varphi_{\tau,1},$$

where I defined $\bar{h} = \frac{h}{1 - \beta \gamma_\tau^x}$, and, using equation (G.8), the evaluation of

$$\varphi_x^\tau + h \varphi_{\tau,1} = \left(\frac{\beta \gamma_\tau^x}{1 - \beta \gamma_\tau^x} T + 1 \right) h \varphi_{\tau,1} = \frac{1 + \beta \gamma_\tau^x (T - 1)}{1 - \beta \gamma_\tau^x} h \varphi_{\tau,1} = \left(1 + \beta \gamma_\tau^x (T - 1) \right) \bar{h} \varphi_{\tau,1}.$$

Using the definition $F \equiv \alpha c \frac{h}{1 - \beta \gamma_\tau^x} \varphi_{\tau,1} = \alpha \bar{c} \bar{h} \varphi_{\tau,1}$ I define the expression

$$\begin{aligned}
\theta_\tau^* &\equiv \beta \gamma_\tau^x (T - 1) = \beta \gamma_\tau^x \left(\frac{2}{1 - F + \sqrt{(1 - F)^2 - 4F \frac{\beta \gamma_\tau^x}{1 - \beta \gamma_\tau^x}}} - 1 \right) \\
&= \beta \gamma_\tau^x \frac{1 + F - \sqrt{(1 - F)^2 - 4F \frac{\beta \gamma_\tau^x}{1 - \beta \gamma_\tau^x}}}{1 - F + \sqrt{(1 - F)^2 - 4F \frac{\beta \gamma_\tau^x}{1 - \beta \gamma_\tau^x}}} \approx \frac{\beta \gamma_\tau^x F}{1 - \beta \gamma_\tau^x - F} \tag{G.16}
\end{aligned}$$

Using this definition and denoting the shadow value of atmospheric carbon under certainty,

see equation (E.5), by $\varphi_{M,1}^{det}$ the term \hat{c} becomes

$$\begin{aligned}
\hat{c} &= \frac{\beta\sigma^{forc}\varphi_{\tau,1}}{M_{pre}} [(\mathbf{I} - \beta\mathbf{\Phi})^{-1}]_{1,1} \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(-1 - \frac{\log(1 - \alpha c(1 + \beta\gamma_\tau^x(T-1))\bar{h}\varphi_{\tau,1})}{\alpha\bar{h}\varphi_{\tau,1}} + \frac{\epsilon(c)}{\delta_\tau} \right) \right) \\
&= \varphi_{M,1}^{det} \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(\frac{-\log(1 - \alpha\bar{h}\varphi_{\tau,1}(1 + \theta_\tau^*))}{\alpha\bar{h}\varphi_{\tau,1}} - 1 + \frac{\epsilon(c)}{\delta_\tau} \right) \right) \\
&= \varphi_{M,1}^{det} \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(\frac{-\log(1 - F(1 + \theta_\tau^*))}{F} - 1 + \frac{\epsilon(c)}{\delta_\tau} \right) \right). \tag{G.17}
\end{aligned}$$

The joined first order approximation in θ_τ^* and F (first approximation), and a first order approximation in F using the definition of θ_τ^* (second approximation) deliver

$$\theta_\tau \equiv \frac{-\log(1 - F(1 + \theta_\tau^*))}{F} - 1 \approx \theta_\tau^* + \frac{1}{2}F \approx \frac{1}{2} \frac{1 + \beta\gamma_\tau^x}{1 - \beta\gamma_\tau^x} F. \tag{G.18}$$

Summarizing the results

Defining $\theta_M^* = \hat{a}\hat{c}$, equation (G.13) delivers

$$\varphi_{M,1} = \hat{c}(1 + \theta_M) \text{ with } \theta_M = \frac{1 - \sqrt{1 - 4\theta_M^*}}{1 + \sqrt{1 - 4\theta_M^*}} \approx \frac{\theta_M^*}{1 - \theta_M^*}.$$

Moreover, equations (G.17) and (G.18) imply

$$\hat{c} = \varphi_{M,1}^{det} \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(\theta_\tau + \frac{\epsilon(c)}{\delta_\tau} \right) \right) \text{ with } \theta_\tau \text{ as in equation (G.18)}. \tag{G.19}$$

Thus, $\varphi_{M,1} = \varphi_{M,1}^{det}(1 + \theta_M) \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(\theta_\tau + \frac{\epsilon(c)}{\delta_\tau} \right) \right)$ and transformed to consumption units

$$SCC = SCC^{det}(1 + \theta_M) \left(1 + \frac{\delta_\tau \bar{h}}{\sigma^{forc}} \left(\theta_\tau + \frac{\epsilon(c)}{\delta_\tau} \right) \right)$$

Finally, equation (G.16) defines θ_τ^* , the expressions $\bar{h} = \frac{h}{1 - \beta\gamma_\tau^x}$ and $F = \alpha\bar{h}\varphi_{\tau,1}$ are as defined in the text above, and equations (G.14), (G.15), and (G.19) give the expression stated for $\theta_M^* = \hat{a}\hat{c}$ in the proposition.

Eliminating δ_τ

The parameter δ_τ can be normalized away as follows. Let the superindex ^{old} denote quantities in the general system above, which I will scale by δ_τ as follows to obtain a rescaled set of equations of motions:

$$h = h^{old}\delta_\tau, c = \frac{c^{old}}{\delta_\tau}, x_t^\tau = \frac{x_t^{old}}{\delta_\tau}, z_t = \frac{z_t^{old}}{\delta_\tau} \text{ and } \epsilon(c) = \frac{\epsilon^{old}(c)}{\delta_\tau}.$$

Generalized temperature follows the equation of motion

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \left(\sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} + h(x_{t+1}^\tau - z_{t+1}) \right) \mathbf{e}_1 \quad (\text{G.20})$$

and hx_{t+1}^τ and hz_{t+1} remain unaffected by the changes. As a consequence, also $\boldsymbol{\tau}_{t+1}$ remains unaffected. Moreover, changing x_t^τ and z_t at all times indeed leads to the following updating equations that no longer depend on δ_τ (and correspond to $\delta_\tau = 1$)

$$\begin{aligned} \mathbb{E} x_{t+1}^\tau &= \mathbb{E} \frac{x_{t+1}^{\tau \text{ old}}}{\delta_\tau} = \gamma_\tau^x \frac{x_t^{\tau \text{ old}}}{\delta_\tau} + \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) = \gamma_\tau^x x_t^\tau + \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \\ \text{Var} x_{t+1}^\tau &= \text{Var} \frac{x_{t+1}^{\tau \text{ old}}}{\delta_\tau} = \frac{1}{\delta_\tau^2} \text{Var} x_{t+1}^{\tau \text{ old}} = \frac{c}{\delta_\tau} \left[2\gamma_\tau^x \frac{x_t^\tau}{\delta_\tau} + \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \right] \\ z_{t+1} &= \frac{z_{t+1}^{\text{old}}}{\delta_\tau} = \gamma_\tau^x \frac{z_t^{\text{old}}}{\delta_\tau} + \left(1 - \frac{\epsilon^{\text{old}}(c)}{\delta_\tau} \right) \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) = \gamma_\tau^x z_t + (1 - \epsilon(c)) \left(\frac{M_{1,t} + G_t}{M_{pre}} - \eta_\tau \right) \end{aligned}$$

Thus, we can set $\delta_\tau = 1$ without loss of generality.

H Welfare, Uncertainty, and Learning

H.1 Proof of Corollary 2

In the following, the terms A , B , and a_{det} refer to those defined on page 55 in equation (E.8) in the proof of the deterministic case. In the present model with uncertainty, I denote the affine shadow value by φ_t^{unc} . I start out with uncertainty only affecting state j , i.e., $\epsilon_t^i = \nu_t^i = 0 \forall i \neq j$. Then, the Bellman equation is

$$\begin{aligned} \sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \kappa_{i,t}^j + \varphi_t^{unc} + B &= A + a_{det} + \beta \varphi_{t+1}^{unc} + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\right. \right. \\ &\quad \left. \left. \alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j (\gamma_i^j \kappa_{i,t}^j + \chi_{i,t}^j) + \varphi_j \epsilon_t^j (\kappa_1^j, \kappa_2^j, \dots) + \varphi_j \nu_t^j \right) \right] \right) \\ \Rightarrow \varphi_t^{unc} + B &= A + a_{det} + \beta \varphi_{t+1}^{unc} + \sum_{i=1}^{\infty} \left[\varphi_{\kappa,i}^j (\beta \gamma_i^j - 1) + \beta \frac{(\alpha \varphi_j)^i}{i! \alpha} \right] \kappa_{i,t}^j \\ &\quad + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\varphi_j \nu_t^j + \sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \chi_{i,t}^j \right) \right] \right). \end{aligned}$$

Matching the coefficients of the new states $\kappa_i^j, i \in \mathbb{N}$, eliminates the squared bracket in front of the cumulants and delivers the shadow values stated in equation (29).

Under certainty, equation (E.8) shows that $\varphi_t^{det} + B = A + a_{det} + \beta \varphi_{t+1}^{det}$. Thus, the difference between the affine shadow value under uncertainty and under certainty in the

present ($t = 0$) is

$$\begin{aligned}\varphi_0^{unc} - \varphi_0^{det} &= \beta(\varphi_1^{unc} - \varphi_1^{det}) + \frac{\beta}{\alpha} \log \left(\mathbb{E}_0 \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \chi_{i,0}^j \right) \right] \right) + G_{\nu_0^j}(\alpha \varphi_j) \\ &= \sum_{t=0}^{\infty} \left[\frac{\beta^{t+1}}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \chi_{i,t}^j \right) \right] \right) + G_{\nu_t^j}(\alpha \varphi_j) \right] \\ &\quad + \lim_{t \rightarrow \infty} \beta^t (\varphi_t^{unc} - \varphi_t^{det}),\end{aligned}$$

where the limit goes to zero for a well-defined model. The welfare difference between the uncertain and the deterministic scenario is characterized by the novel cumulant dependent part of welfare $\sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,t}$ and the affine contribution to welfare $\varphi_0^{unc} - \varphi_0^{det}$

$$\begin{aligned}\Delta W^j &= V_0^{unc} - V_0^{det} = \sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \kappa_{i,0}^j + \varphi_0^{unc} - \varphi_0^{det} \\ &= \sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \kappa_{i,0}^j + \sum_{t=0}^{\infty} \frac{\beta^{t+1}}{\alpha} G_{\nu_t^j}(\alpha \varphi_j) \\ &\quad + \sum_{t=0}^{\infty} \frac{\beta^{t+1}}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i}^j \chi_{i,t}^j \right) \right] \right)\end{aligned}$$

For the general case of shocks to more than one state j , it is easy to see the additivity of the individual contributions for each j giving rise to the total welfare loss ΔW^{unc} stated in equation (??) in Proposition 2.

H.2 Proof of Corollary 1

In the case of persistent carbon sink shocks, the adjustments in the equations of motion (22) and (??) modify or add the following terms to the Bellman equation (9)³²

$$\varphi_{\epsilon} \epsilon_t + \varphi_t + \dots = \dots + \beta \varphi_{t+1} + \beta \varphi_{\epsilon} \gamma \epsilon_t + \beta [\varphi_{M_1} - \varphi_{M_2}] \epsilon_t + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp [\alpha \varphi_{\epsilon} \chi_t] \right).$$

It is easily observed that these changes do not affect the optimal investment rate and labor distribution. Matching the coefficients of the flow adjustment ϵ_t to make the Bellman equation independent of its level delivers equation (??) for the shadow value φ_{ϵ} . The remaining terms imply $\varphi_t = \beta \varphi_{t+1} + \frac{1}{\alpha} \log \left(\mathbb{E}_t \exp [\alpha \beta \varphi_{\epsilon} \chi_t] \right) + const_t$, where $const_t$ is a term that is independent of the uncertainty. Given $\epsilon_0 = 0$, the welfare difference between the deterministic and the uncertain scenario is determined by the difference of the affine value function contributions

$$\begin{aligned}\Delta W^{VAR} &= V_0^{unc} - V_0^{det} = \varphi_0^{unc} - \varphi_0^{det} = \beta(\varphi_1^{unc} - \varphi_1^{det}) + \frac{\beta}{\alpha} \log \left(\mathbb{E}_0 \exp [\alpha \varphi_{\epsilon} \chi_0] \right) \\ &= \sum_{i=0}^{\infty} \frac{\beta^{i+1}}{\alpha} \log \left(\mathbb{E}_i \exp [\alpha \varphi_{\epsilon} \chi_i] \right) + \lim_{i \rightarrow \infty} \beta^i (\varphi_i^{unc} - \varphi_i^{det}).\end{aligned}$$

³²Under uncertainty, let Ω denote the underlying Borel sigma algebra and \mathcal{F} the filtration generated by the stochastic process. The equations of motion are conditional on $\omega \in \Omega$, the controls are adapted to the filtration, and expectation in period t are conditional on \mathcal{F}_t .

For a well-defined dynamic system $\lim_{i \rightarrow \infty} \beta^i (\varphi_{t+i}^{unc} - \varphi_{t+i}^{det}) = 0$ and I obtain the general welfare loss equation for non-stationary shocks

$$\Delta W^{VAR} = \frac{1}{\alpha} \sum_{t=0}^{\infty} \beta^{t+1} \log [\mathbb{E} \exp [\alpha \varphi_{\epsilon} \chi_t]] \quad . \quad (\text{H.1})$$

For a sequence of identically distributed shocks χ_t , I obtain the welfare cost of uncertainty stated in (30) by evaluating the implied geometric sum in equation (H.1).

H.3 Proof of Proposition ??

In the case of anticipated learning, the new equation of motion for the atmospheric and the biosphere-and-upper-ocean carbon reservoirs take the form

$$M_{1,t+1} = (\Phi \mathbf{M}_t)_1 + \sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} + \epsilon_t + \nu_t, \quad (\text{H.2})$$

$$M_{2,t+1} = (\Phi \mathbf{M}_t)_2 - \epsilon_t - \nu_t.$$

I model the learning process based on atmospheric carbon observation.³³ Rearranging equation (H.2), the decision maker derives information on ϵ_t from the realizations

$$\hat{\epsilon}_t = M_{1,t+1} - (\Phi \mathbf{M}_t)_1 - \sum_{i=1}^{I^d} E_{i,t} - E_t^{exo} - \nu_t \quad .$$

The equations of motion for the Bayesian prior's mean and variance are

$$\mu_{\epsilon,t+1} = \frac{\sigma_{\epsilon,t}^2 \hat{\epsilon}_t + \sigma_{\nu,t}^2 \mu_{\epsilon,t}}{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2} \quad \text{and} \quad \sigma_{\epsilon,t+1}^2 = \frac{\sigma_{\nu,t}^2 \sigma_{\epsilon,t}^2}{\sigma_{\nu,t}^2 + \sigma_{\epsilon,t}^2} \quad .$$

This standard Bayesian updating equation characterizes the posterior mean as a weighted average of the new observation and its prior mean. The weight of the new observation is inversely proportional to the variance of the measurement error (or proportional to its precision). The weight on the prior's mean is inversely proportional to its variance. The variance of the carbon cycle uncertainty in this Bayesian learning model falls exogenously over time. The smaller the ratio of stochasticity to overall uncertainty $\frac{\sigma_{\nu,t}^2}{\sigma_{\nu,t}^2 + \sigma_{\epsilon,t}^2}$, the faster the learning.

³³In principle, the decision-maker could simultaneously learn from observing the carbon concentration in the combined biosphere and upper ocean reservoir. However, whereas the CO₂ concentration in the atmosphere is somewhat homogenous, the concentration (partial pressure) in the ocean varies from 250 ppm to 500ppm over regions and seasons, and an annual 1Gt ocean uptake is driven by as little as a 2ppm difference between the concentrations in the atmosphere and the oceans. Thus, measurement errors in the non-atmospheric carbon reservoir are so much larger that an observation-based learning model can comfortably ignore these additional measurements.

These adjustments in the equations of motion imply modifications of the Bellman equation (9) captured by the terms

$$\begin{aligned} \varphi_\mu \mu_{\epsilon,t} + \varphi_t + \dots &= \dots + \beta \varphi_{t+1} + \beta \varphi_\mu \frac{\sigma_{\nu,t}^2}{\sigma_{\nu,t+1}^2 + \sigma_{\epsilon,t}^2} \mu_{\epsilon,t} \\ &+ \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\varphi_{M_1} - \varphi_{M_2} + \varphi_{\mu,t} \frac{\sigma_{\epsilon,t}^2}{\sigma_{\nu,t+1}^2 + \sigma_{\epsilon,t}^2} \right) (\epsilon_t + \nu_t) \right] \right). \end{aligned} \quad (\text{H.3})$$

Matching the coefficients of the informational state $\mu_{\epsilon,t}$ to make the Bellman equation independent of its level delivers equation (??) for the shadow value φ_μ . Solving inductively the remaining state-independent terms in equation (H.3) for the welfare difference between the uncertain and the deterministic scenario as in the proof of Proposition 1 delivers the welfare loss

$$\begin{aligned} \Delta W^{Bayes} &= \sum_{t=0}^{\infty} \beta^{t+1} \alpha \frac{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2}{2} \left[\varphi_{M_1} - \varphi_{M_2} + \varphi_{\mu,t} \frac{\sigma_{\epsilon,t}^2}{\sigma_{\nu,t}^2 + \sigma_{\epsilon,t}^2} \right]^2 \\ &= \sum_{t=0}^{\infty} \beta^{t+1} \alpha \frac{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2}{2} (\varphi_{M_1} - \varphi_{M_2})^2 \left[\frac{(1-\beta)\sigma_{\epsilon,t}^2 + (1-\beta)\sigma_{\nu,t}^2 + \beta\sigma_{\epsilon,t}^2}{(1-\beta)(\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2)} \right]^2, \end{aligned}$$

where I inserted the shadow value φ_μ from equation (??). Canceling terms in the numerator of the expression in squared brackets delivers equation (27) in the main text.

H.4 Quantitative Analysis of Carbon Cycle Uncertainty

The quantification of carbon cycle uncertainty in section ?? is an informed guess based on Joos et al.'s (2013) model comparison study and the measurement error implied by the missing sink. Here, I attempt to bound the welfare impact using a somewhat reasonable lower and upper bound for carbon cycle uncertainty. In the VAR model of section ??, the left panel of Figure 7 reduces the shock's standard deviation to 10 Gt per decade. It builds up to a 200Gt standard deviation after about 300 years, which is significantly lower than the 500Gt standard deviation in Joos et al.'s (2013) model comparison study. The resulting welfare loss is approximately 28 billion USD. The right panel of Figure 7 increases the shock's standard deviation to 50 Gt per decade. It builds up to the suggested 500Gt standard deviation after 125 years, but implies double that value after around 350 years. The resulting welfare loss is approximately 700 billion USD.

In section ??, I found a willingness to pay for a stochasticity reduction (or reduction in measurement error) of approximately half a billion USD per Gt decadal standard deviation. If the initial measurement error σ_ν is already down to 5Gt instead of 10Gt per decade, then this willingness to pay is also cut into half to approximately 260 million USD. If the initial measurement error is doubled ($\sigma_\nu = 20\text{Gt}$), then the willingness to pay increases to 750 million USD.

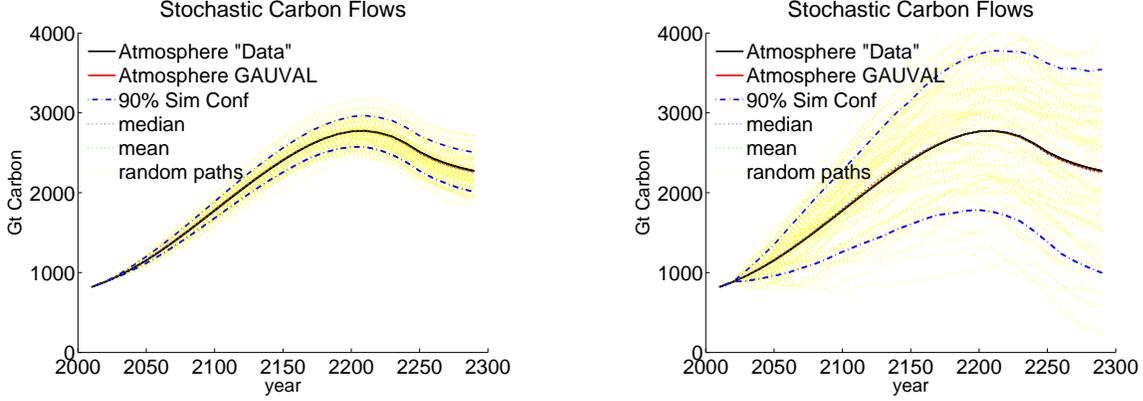


Figure 7: shows the evolution of atmospheric carbon under the low and the high specifications of the carbon cycle shock in equation (??), $\sigma_\chi = 10$ Gt on the left and $\sigma_\chi = 50$ Gt on the right. The shock's persistence of $\gamma_M = 0.997$ is calibrated to Joos et al.'s (2013) model comparison study. The underlying emission scenario is DICE's business as usual. The deterministic DICE evolution (5 year time steps, "Data"), the deterministic GAUVAL evolution (10 year time steps), and the mean and the median of 1000 uncertain trajectories are hardly distinguishable.

H.5 Proof or Proposition ??

In the combined model of persistent epistemological and VAR uncertainty over the temperature increase in section ?? the Bellman equation gains the following terms

$$\begin{aligned} \sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,t} + \varphi_t + \dots &= \dots + \beta \varphi_{t+1} + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i} (\gamma_i \kappa_{i,t} + \chi_{i,t}^\tau) \right. \right. \right. \\ &\quad \left. \left. \left. + \varphi_{\tau,1} \epsilon_t^\tau (\kappa_1, \kappa_2, \dots) \right) \right] \right). \\ \Rightarrow \varphi_t + \dots &= \dots + \beta \varphi_{t+1} + \beta \sum_{i=1}^{\infty} \left[\varphi_{\kappa,i} (\beta \gamma_i - 1) + \beta \frac{(\alpha \varphi_{\tau,1})^i}{i! \alpha} \right] \kappa_{i,t} \\ &\quad + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i} \chi_{i,t}^\tau \right) \right] \right). \end{aligned}$$

Matching the coefficients of the new states κ_i , $i \in \mathbb{N}$, eliminates the squared bracket in front of the cumulants and delivers the shadow values stated in equation (29). The difference between the uncertain and the deterministic value function's affine components derives analogously to the proof of Proposition 1 to

$$\begin{aligned} \varphi_0^{unc} - \varphi_0^{det} &= \beta (\varphi_1^{unc} - \varphi_1^{det}) + \frac{\beta}{\alpha} \log \left(\mathbb{E}_0 \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i} \chi_{i,t}^\tau \right) \right] \right) \\ &= \sum_{t=0}^{\infty} \frac{\beta^{t+1}}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i} \chi_{i,t}^\tau \right) \right] \right). \end{aligned}$$

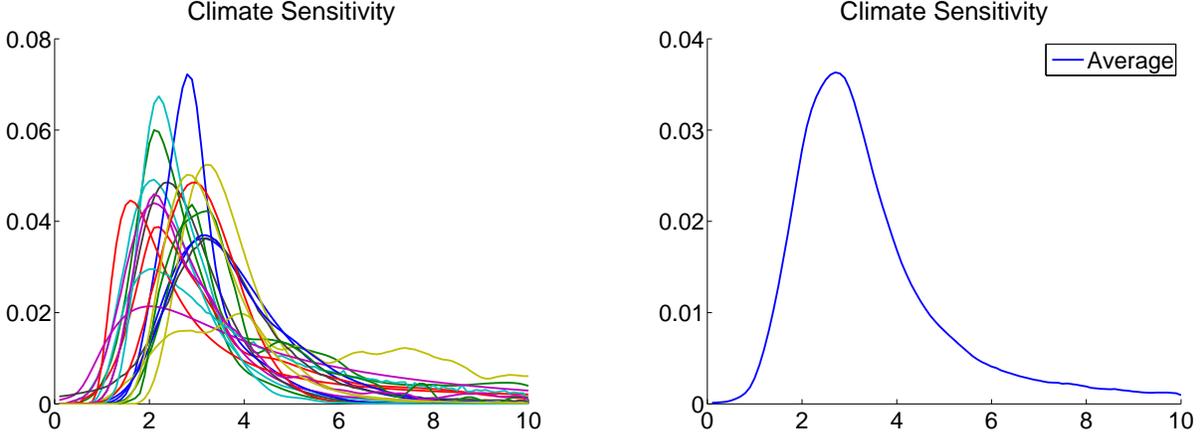


Figure 8: shows estimates of the probability distribution of a temperature increase in degree Celsius resulting from a doubling of CO₂ concentrations with respect to industrial levels. On the left, the figure depicts 20 probability distributions of climate sensitivity derived by different groups and using different methodological approaches (Meinshausen et al. 2009). On the right, the figure depicts the average distribution assigning equal weight to each approach. The Figure is to be interpreted as the probability density of a temperature increase conditional on not exceeding 10C.

The welfare difference between the uncertain and the deterministic scenario is now comprised of a state (cumulant) dependent part $\sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,t}$ and the affine part of the value functions

$$\begin{aligned} \Delta W^{temp} &= V_0^{unc} - V_0^{det} = \sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,t} + \varphi_0^{unc} - \varphi_0^{det} \\ &= \sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,t} + \sum_{i=0}^{\infty} \frac{\beta^{i+1}}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\sum_{i=1}^{\infty} \varphi_{\kappa,i} \chi_{i,t}^{\tau} \right) \right] \right). \end{aligned}$$

In the case of identically distributed shocks over time, the second sum characterizes a geometric series giving rise to the factor $\frac{\beta}{1-\beta}$, turning the welfare loss into the form stated in equation (28) in the main text.

H.6 Quantitative Analysis of Temperature Uncertainty

Figure 8 illustrates the uncertainty governing the temperature increase from a doubling of the CO₂ concentration, the so-called climate sensitivity. On the left, the figure depicts 20 probability distributions of climate sensitivity derived by different groups and using different methodological approaches (Meinshausen et al. 2009). These probability densities are conditional on the temperature increase not exceeding 10 C. On the right, the figure depicts the average distribution assigning equal weight to each approach. The distribution is positively skewed and exhibits more weight in the right tail as compared to a (truncated) normal distribution. It serves as the starting point for my numeric estimates of the welfare loss from temperature uncertainty.

This average climate sensitivity distribution on the right of Figure 8 has an expected value of 3.4 C, differing from the common best guess of 3 C employed so far. Focusing on

the uncertainty contribution, I shift Meinshausen et al.’s (2009) distribution to the left to conserve the 3 C warming expectation. I denote the implied distribution of the generalized temperature state by $\tilde{\tau}^\infty$. By equation (??), the temperature flow uncertainty $\epsilon^\tau = [\mathbf{1} - \boldsymbol{\sigma}]_{1,1} \tilde{\tau}^\infty - 2\sigma^{forc}$ generates this long-run temperature uncertainty under the assumption of a doubling of preindustrial CO₂ concentrations. I start by assuming only the VAR model, which corresponds to autoregressive shocks χ_1 to the mean. Such shocks build up over time, and for a doubling of CO₂ concentrations a stationary shock $\chi_1 = (1 - \gamma)\epsilon^\tau$ generates the depicted distribution of climate sensitivity. As explained in section ??, the simulation assigns a fraction ζ of the long-run climate sensitivity uncertainty to this shock-based contribution, and the fraction $1 - \zeta$ to the initial epistemological uncertainty. More than two decades of IPCC assessment reports have not tightened the confidence interval on climate sensitivity. Therefore, I assume a persistence of epistemological (and VAR shock) uncertainty of $\gamma = 0.9$ in my “baseline” scenario. In evaluating the welfare loss from temperature uncertainty along the DICE business as usual scenario, I scale the exogenous shocks $\chi_{1,t}$ proportional to the atmospheric CO₂ concentrations along the business as usual path (thick black ‘data’ line in Figure 4).³⁴

³⁴The scaling of the shock is proportional to the CO₂ concentration because the shock affects transformed temperature, which translates logarithmically into real temperature, accounting for falling radiative forcing from an additional ton of CO₂.