



Working Papers

www.cesifo.org/wp

Closed-Form Integrated Assessment and Uncertainty

Christian P. Traeger

CESIFO WORKING PAPER NO. 5464
CATEGORY 10: ENERGY AND CLIMATE ECONOMICS
AUGUST 2015

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

Closed-Form Integrated Assessment and Uncertainty

Abstract

The paper derives the optimal carbon tax in closed-form from an integrated assessment of climate change. The formula shows how carbon, temperature, and economic dynamics quantify the optimal mitigation effort. The model's descriptive power is comparable to numeric models used in policy advising. Uncertainty surrounding climate change remains large, and the paper derives closed-form expressions of welfare loss from shocks and epistemological uncertainty. These expressions interact (intrinsic) risk attitude, distributional moments, and the climatic shadow values, and exhibit different sensitivities to time preference. Welfare gains from reducing uncertainty about temperature feedbacks are much higher than the gains from better measurements of carbon flows.

JEL-Code: Q540, H430, E130, D810, D900, D610.

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

Christian P. Traeger
Department of Agricultural & Resource Economics
University of California Berkeley
207 Giannini Hall #3310
USA - 94720-3310 Berkeley CA
traeger@berkeley.edu

August 2015

I am grateful to Larry Karp, Till Requate, Sverre Jensen, Fortunat Joos, Malte Meinshausen, and the seminar participants at the Universities of Oslo, Berkeley, Stanford, Miami, Toulouse, Cambridge, Amsterdam, Ohio State, and at the meetings of the Verein fuer Socialpolitik, the Association of Environmental and Resource Economists, and the Association for Public Economic Theory. This research is supported by the National Science Foundation through the Network for Sustainable Climate Risk Management GEO-1240507.

1 Introduction

How strictly should we regulate carbon emissions? Quantitative policy advice relies entirely on complex numeric models. This paper provides a closed-form expression that shows how the different model components interact to determine the optimal carbon tax. It also provides closed-form expressions of the welfare costs arising from the major climate change uncertainties. These formulas help to identify research priorities.

The closed-form results deepen insights, make it easier to identify and understand sensitivities, and make analytic proofs possible. The model overcomes Bellman’s curse of dimensionality, which limits integrated assessment under uncertainty. By imbedding a high-dimensional science-based calibration in a tractable framework, the model provides quantitative policy advice together with improved qualitative understanding.

The present model is intended to accompany, not to replace, numeric approaches to the integrated assessment of climate change. Yet, it has (at least) similar descriptive power as DICE, the most widely used numeric integrated assessment model. It contains a more general energy sector and a slightly better fit of the temperature dynamics. Moreover, by disentangling risk aversion from intertemporal consumption smoothing it improves the calibration in ways important to both risk assessment and discounting. Compared to DICE, the model has a simplified capital decay structure.

Environmental economists have a long history of using closed-form models to examine qualitative aspects of climate change policy. Recently, Golosov et al. (2014) broke new ground in developing and calibrating an analytically tractable IAM. Their model combines a Ramsey growth model with an explicit model of the energy sectors and a linear impulse response of economic production to carbon emissions. Gerlagh & Liski (2012) extend the model to account for the empirically important delay between atmospheric CO₂ concentrations and economic impact. 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), and intergenerational games (Karp 2013). In tribute to the pioneering work by Golosov, Hassler, Krusell, Tsyvinski, Gerlagh, and Liski, I call the model in this paper GAUVAL, taking the *i*th letter of the *i*th author.

GAUVAL marries Golosov et al.’s (2014) base model with a standard climate model that shares the same structure as all numeric IAMs used in policy advising. This climate model consists of (i) a standard carbon cycle model, (ii) the physical non-linearity in the response of radiative forcing to atmospheric CO₂ accumulation (“greenhouse effect”), and (iii) a model of the atmosphere-ocean temperature dynamics. Components (ii) and (iii) are new with respect to its analytic predecessors.¹ I derive the general conditions under which

¹Closest in the analytic literature are Gerlagh & Liski (2012) who jointly approximate component (ii), component (iii), and the translation back into damages by a delay parameter interpreted as warming delay, and Anderson et al. (2014) who employ the assumption that temperatures are approximately linear in cumulative post-industrial emissions. The present model also follows the numeric models in adding (exogenous)

such a model solves in closed form. The explicit climate change model permits a direct component-wise comparison to standard IAMs and to scientific climate change models, and it permits a direct calibration of the temperature dynamics to the IPCC (2013) scenarios. Most importantly, the climate model allows me to identify and compare the roles of carbon and temperature dynamics in the optimal carbon tax and in the welfare loss. The explicit tracking of temperature dynamics is necessary to compare the implications of carbon cycle and temperature response uncertainty, the two main scientific climate change uncertainties.

Despite decades of research, uncertainty remains one of the most widely accepted and dreaded characteristics of climate change. The present model characterizes the welfare loss from uncertainty as a function of the distribution's moments and corresponding powers of intertemporal risk aversion. I examine the consequences of persistent shocks, epistemological uncertainty, and anticipated learning. Kelly & Kolstad (1999) pioneer the explicit treatment of uncertainty in numeric IAMs. Recent extensions disentangle risk aversion from intertemporal consumption smoothing (Crost & Traeger 2014, Jensen & Traeger 2014), and examine the relevance of thick tails (Kelly & Tan 2013). The present research takes examines these issues in an analytic framework, and compares uncertainty governing the carbon flows to the uncertainty that governs the temperature response to a given atmospheric carbon concentration (climate sensitivity).

Golosov et al.'s (2014) original work analyzes damage uncertainty and is significantly extended by Gerlagh & Liski (2014). Welfare effects in these models depend on the expectation but not the higher moments of random or uncertain parameters. I show that this strong certainty equivalence results from intertemporal risk neutrality in the case of entangled logarithmic Arrow Pratt risk aversion. Li et al. (2014) and Anderson et al. (2014) use a version² of Golosov et al.'s (2014) model to show that robust preferences can have a large, and qualitatively ambiguous effect on the optimal carbon tax.

2 The Model

GAUVAL'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, discusses the underlying calibration, and introduces preferences that disentangle risk aversion from

CO₂ emissions from land use change and forestry and accounting for the (exogenous) radiative forcing from other greenhouse gases.

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.

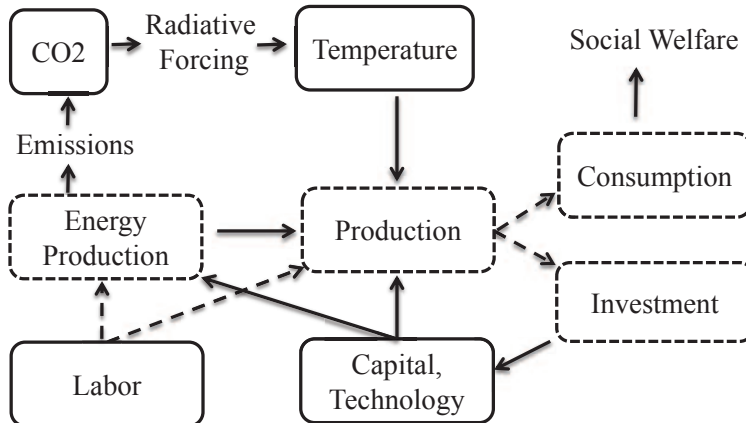


Figure 1: The structure of GAUVAL and most Integrated Assessment Models. Solid boxes characterize the model’s state variables, dashed boxes are flows, and dashed arrows mark choice variables.

intertemporal consumption smoothing.

2.1 The Economy

Utility is logarithmic and the social planner’s time horizon is infinite. The logarithmic utility function captures only (deterministic) 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

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

The aggregate energy input E_t is a smooth and monotonic function

$$E_t = g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)) \tag{1}$$

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 CO_2 (“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 CO_2 (“clean”). I assume a system of energy sectors of the general form (1) 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 $\mathbf{R}_0 \in \mathbb{R}_+^{I^d}$ 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, $D_t(0) = 0$. Proposition 1 characterizes the class of damage functions $D_t(T_{1,t})$ that permit an analytic solution of the model.

Following Golosov et al.'s (2014) assumption of full depreciation after one period, the capital stock's equation of motion is

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})] - C_t. \quad (2)$$

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. Appendix A extends the model to allow for more sophisticated capital persistence by interacting an exogenous capital growth rate approximation with capital depreciation. This step enables GAUVAL to match the empirical capital accumulation and it makes the social planner aware of the additional investment payoff from higher capital persistence. The extension does not affect the equations describing the optimal carbon tax or the welfare equations. The crucial implication of equation (2), and of the empirically better founded extension in Appendix A, is that the investment rate will be independent of the system state. The investment rate is approximately independent of the climate states also in DICE (see Appendix A). The (remaining) channels for climate policy are the restructuring of the energy sector, investment levels, and a shift of labor allocation between the final good sector and the energy sectors.

2.2 The Climate System

The energy sector's CO₂ emissions enter the atmosphere. We also emit (smaller quantities of) CO₂ through land conversion, forestry, and agriculture. 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.

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.

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 reservoir and the matrix Φ capture the transfer coefficients. Then

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

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

$$F_t = \eta \frac{\log \frac{M_{1,t} + G_t}{M_{pre}}}{\log 2} . \quad (4)$$

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 only with delay. After several centuries, the new equilibrium⁴ 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 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 $T_{0,t} = \frac{s}{\eta} F_t$. A given ocean layer slowly adjusts its own temperature to the temperature of the surrounding layers. I model next

⁴The 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.

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 $T_{i+1,t}$.

A 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.⁵ I collect all weights in the $l \times l$ matrix $\boldsymbol{\sigma}$, which characterizes the heat exchange between the atmosphere and the different ocean layers.

Next period's temperature in layer i is the generalized mean of its own and the adjacent present period temperatures

$$T_{i,t+1} = \mathfrak{M}_i^{\boldsymbol{\sigma}}(T_{i,t}, w_i^{-1}T_{i-1,t}, w_{i+1}T_{i+1,t}) \text{ for } i \in \{1, \dots, l\}. \quad (5)$$

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 empirical calibration in section 2.4 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.

2.3 Solving GAUVAL

Appendix B solves GAUVAL 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 motions that are linear in the state variable, and on control variables that are additively separable from the states. GAUVAL 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 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 B.

⁵For notational convenience equation (5) below will write a mean of three temperature values also for the deepest layer ($i = l$) with a zero weight on the arbitrary entry T_{l+1} .

Proposition 1 *An affine value function of the form*

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t \quad (6)$$

solves GAUVAL 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 takes the form

$$D(T_{1,t}) = 1 - \exp[-\xi_0 \exp[\xi_1 T_{1,t}] + \xi_0], \quad \xi_0 \in \mathbb{R},$$

the mean in the equation of motion (5) for temperature layer $i \in \{1, \dots, l\}$ takes the form

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

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 $w_i, i \in \{1, \dots, l-1\}$).

The coefficients φ in the value function are the shadow value 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 2.4 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

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1. \quad (8)$$

The parameter σ^{forc} is the weight on radiative forcing in the atmospheric temperature's equation of motion. It determines the speed at which temperatures respond to the greenhouse effect. To achieve additive separability between controls and states, the consumption rate $x_t = \frac{C_t}{Y_t[1-D_t(T_t)]}$ replaces absolute consumption as the consumption-investment control. Under the assumptions of Proposition 1, the optimal consumption rate is

$$x_t^* = 1 - \beta \kappa. \quad (9)$$

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.

2.4 Calibration

I employ the carbon cycle of DICE 2013. Running the model in a 10 year time step, I double the original 5 year transition coefficients. Figure 4 in section 4.2 confirms that the rescaled 10 year transition matrix implies an evolution of the carbon stock that is indistinguishable from that of the original 5 year step of DICE 2013. I employ the usual capital share $\kappa = 0.3$ and use the International Monetary Fund's (IMF) 2015 investment rate forecast $1 - x^* = 25\%$ to calibrate pure time preference. Equation (9) implies $\beta = \frac{1-x^*}{\kappa} = \frac{0.25}{0.3}$ and an annualized rate of pure time preference of $\rho = \frac{1}{10} \log \beta = 1.75\%$. The conversion of utility values into 2015 USD relies on the log utility's implication that $dC = C du = x Y du$, where the consumption rate is $x = 75\%$ and Y is equal to ten (time step) times the IMF's global economic output forecast of $Y_{2015} = 81.5$ trillion USD.

Economic damage functions are crucial and yet hard to determine. The most widespread IAM DICE uses the form $D(T) = \frac{1}{1+0.0028T^2}$. Nordhaus (2008) calibrates the coefficient 0.0028 based on a damage survey for a 2.5C warming. I calibrate GAUVAL'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 a more convex damage curve passing through Nordhaus' calibration point at 2.5°C is just as likely. GAUVAL's damage function initially generates damages that are slightly higher than in DICE-2007 and matches them exactly at 2.5°C. For higher temperatures up to a 12°C warming, my base case calibration delivers slightly lower damages compared to DICE, and it generates higher damages for global warming above 12°C, warming levels that imply a hard-to-conceive change of life on the planet. Figure 2 also depicts two dashed variations of GAUVAL'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. Section 3 discusses how such changes affect welfare and the optimal carbon tax.

The calibration of temperature dynamics (equation 7) uses the emission scenarios of the recent assessment report of the Intergovernmental Panel on Climate Change IPCC (2013). These so-called Representative Concentration Pathways (RCP) replace the (SRES-) scenarios of the earlier assessment reports (Moss et al. 2007). They are labeled by the approximate radiative forcing levels they produce by the end of the century (measured in W/m^2). These new RCP scenarios are defined for longer time horizons and, thus, better suited for calibration

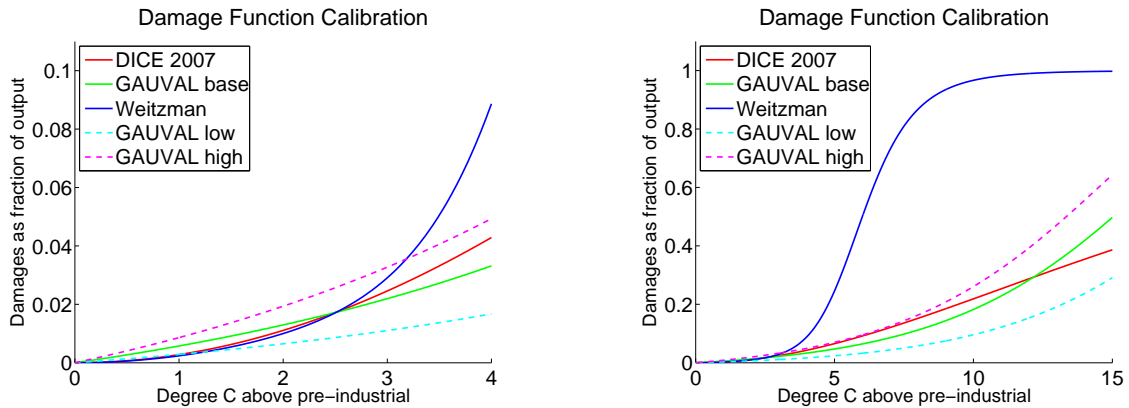


Figure 2: GAUVAL’s damage function compared to that of DICE-2007 and a highly convex damage function suggested by Weitzman (2010). All three lines coincide for a 2.5°C warming, the common calibration point based on Nordhaus (2008). The dashed curves depict GAUVAL’s damage function for a $\pm 50\%$ variation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

than the earlier SRES scenarios. 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 atmospheric ocean general circulation models (AOGCMs) and is employed in the IPCC’s assessment report. DICE was calibrated to one of the old SRES scenarios using an earlier version of Magicc. My calibration of GAUVAL uses three ocean layers (upper, middle, and deep) 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 GAUVAL’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. GAUVAL’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.

2.5 Uncertainty

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

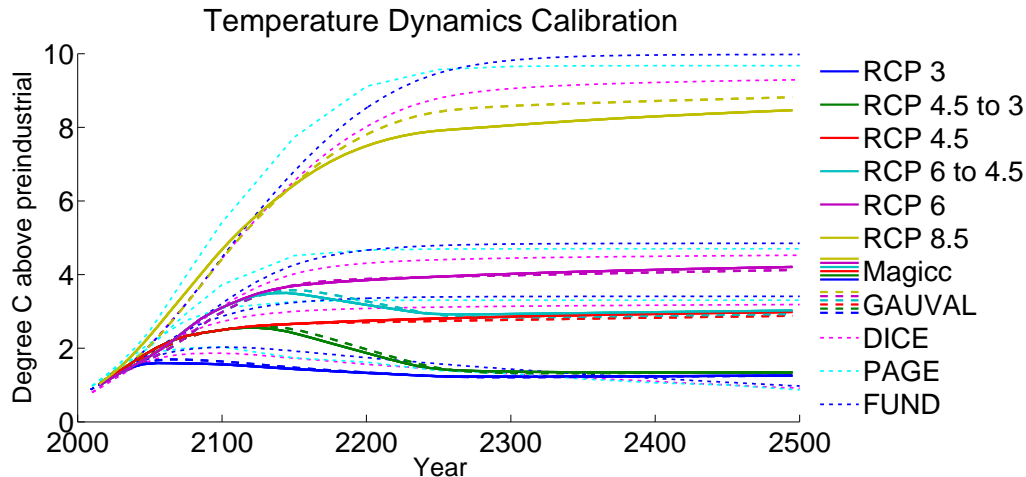


Figure 3: GAUVAL’s response vis a vis Magicc’s response to the color coded radiative forcing scenarios used in the latest IPCC assessment report. RCP 3 is the strongest stabilization scenario and RCP 8.5 is a business as usual scenario. The Magicc model (solid lines) emulates the large scientific models and is used in the IPCC’s assessment reports. GAUVAL (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. GAUVAL performs better in all scenarios. The scenarios RCP 4.5 to 3 and RCP 6 to 4.5 are scenarios switching from a higher to a lower forcing trajectory. They are not part of the IPCC’s official scenario selection but useful to calibrate the model across a temperature peak.

et al. 2010, Chen et al. 2013, Bansal et al. 2012).⁶ Merely increasing the utility function’s curvature would result in a much larger risk-free discount rate than observed in the markets (risk-free rate puzzle). The market rejects the assumption that the intertemporal elasticity of substitution fully determines risk attitude. This assumption is built into the standard intertemporally additive expected utility model and implies a form of risk neutrality in intertemporal choice (Traeger 2014).

I follow the asset pricing literature, an increasing strand of macroeconomic literature, and some recent numeric approaches to climate change assessment (Croft & Traeger 2014, Jensen & Traeger 2014) in using Epstein-Zin-Weil preferences. This approach accommodates a realistic coefficient of risk aversion, disentangling it from the unit elasticity of intertemporal substitution. The extension results in a Bellman equation with a non-linear risk aggregation

$$V(k_t, \tau_t, \mathbf{M}_t, \mathbf{R}_t, t) = \max_{x_t, \mathbf{N}_t} \log c_t + \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(V(k_{t+1}, \tau_{t+1}, \mathbf{M}_{t+1}, \mathbf{R}_{t+1}, t) \right) \right] \right). \quad (10)$$

Expectations \mathbb{E}_t are conditional on time t information. In general, consumption and next

⁶Nakamura 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.

period's states are uncertain. The non-linear uncertainty aggregator is a generalized mean $f^{-1} \mathbb{E}_t f$ with $f(\cdot) = \exp(\alpha \cdot)$.

A negative parameter α characterizes risk aversion in the intertemporal sense, axiomatically defined as intrinsic risk aversion in Traeger (2014). 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. Appendix D explains the relation between equation (10) and Epstein & Zin's (1991) original formulation for this special case of an intertemporal elasticity of substitution of unity. The coefficient of constant relative risk aversion in Epstein & Zin's (1991) definition of Arrow-Pratt risk aversion is $\text{RRA} = 1 - \alpha^* = 1 - \frac{\alpha}{(1-\beta)}$. The asset pricing literature estimates RRA in the range $[6, 9.5]$ corresponding to $\alpha \in [-1, -1.5]$. Note that Epstein-Zin preferences face the same issue as standard expected utility theory when it comes to calibrating risk aversion in the small and in the large (Rabin 2000): calibrating aversion on small bets requires degrees of risk aversion that are unreasonably high for large bets. Consequently, I use $\alpha = -1$ for quantitative examples with high uncertainty and $\alpha = -1.5$ for quantitative examples with low uncertainty. Figure 6 in Appendix D illustrates 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.

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 CO₂ 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 B solves for the shadow values and derives the following result on optimal carbon dioxide taxation.

Proposition 2 *Under the assumptions of section 2.1 (economy) and section 2.2 (climate*

system) the SCC in money-measured consumption equivalents is

$$SCC_t = \frac{\beta Y_t}{M_{pre}} \underbrace{\xi_0}_{2.2\%} \underbrace{[(\mathbf{1} - \beta \boldsymbol{\sigma})^{-1}]_{1,1}}_{1.4} \underbrace{\sigma^{forc}}_{0.42} \underbrace{[(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1}}_{3.7} = 56.5 \text{ \$/tC} \quad (11)$$

$25.5 \frac{\$}{tC}$

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 2.4.

As emphasized by Golosov et al. (2014), the SCC is proportional to production Y_t and increases over time at the rate of economic growth. In the present formula, 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 $\tau_1 = \exp(\xi_1 T_1)$. These terms together would imply a carbon tax of 25.5\$ per ton of carbon.

The subsequent terms paint a detailed picture of the climate dynamics. Appendix C provides a simple illustrations for a two layer carbon cycle and atmosphere-ocean temperature system. A von Neumann series expansion of the (bounded operator) $\beta \boldsymbol{\Phi}$ helps to interpret the general term $[(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1}$ governing carbon cycle dynamics

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

The element $[\boldsymbol{\Phi}^i]_{1,1}$ 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., $[\boldsymbol{\Phi}^2]_{1,1} = \sum_j \Phi_{1,j} \Phi_{j,1}$ 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. In summary, the term $[(\mathbf{1} - \beta \boldsymbol{\Phi})^{-1}]_{1,1}$ 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 when ignoring warming delay and the temperature's atmosphere-ocean interaction.

The terms $[(\mathbf{1} - \beta \boldsymbol{\sigma})^{-1}]_{1,1} \sigma^{forc}$ capture the atmosphere-ocean temperature delay dynamics. Analogously to the interpretation in the case of carbon, the expression $[(\mathbf{1} - \beta \boldsymbol{\sigma})^{-1}]_{1,1}$ characterizes the generalized heat flow that enters, stays, and returns to the atmospheric layer. Note that the simple closed-form expression in equation (11) 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 atmospheric temperature responds to radiative forcing. The response delay, around 0.4, significantly 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.

Expressed in tons of CO₂, this SCC is 15.5 USD, coinciding up to one dollar to the DICE-2007 carbon tax for 2015.⁷ 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 70%. Ignoring carbon persistence would result in a carbon tax that is only 27% of its actual value.

The earlier models of Golosov et al. (2014) and Gerlagh & Liski (2012) use a carbon decay model mostly equivalent to the carbon cycles commonly employed in IAMs and adopted in GAUVAL. The social cost of carbon in these papers contains a summation term that similarly translates carbon persistence into a multiplier of an output loss based contribution. These 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 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 GAUVAL’s factor of $1.4 \cdot 0.42 \approx .6$ based on an explicit model of temperature dynamics.

Embedded in equation (11) 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

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

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

⁷DICE expresses the carbon tax in 2015 in USD of 2005, which have to be translated into 2015 USD for the comparison.

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

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

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.⁸ 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 B.3 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. GAUVAL 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 (GAUVAL does not have an explicit damage function for ocean acidification).

3.2 The Optimal Carbon Tax: A Declaration of Independence

In general, the optimal carbon tax is the SCC along the optimal emission trajectory. The SCC in equation (11) is independent of the absolute stock of carbon in the atmosphere. In consequence, the SCC in GAUVAL is independent of the future emission trajectory, and the SCC directly specifies the optimal carbon tax. This finding already prevails in Golosov et al. (2014). While convenient, it raises a certain degree of discomfort: our optimal effort to reduce a ton of carbon is independent of whether we are in a world of high or low carbon concentrations, and independent of whether we are in a world of high or low prevailing temperatures. This discomfort only increases when we learn that a fraction of any emitted ton of carbon stays in the atmosphere for at least thousands of years. A common argument governing climate change action is: if we delay mitigation today, we have to do even more

⁸Ocean 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.

tomorrow. The model tells us: if we delay policy today, we have to live with the consequences, but we do not have to compensate in our future mitigation effort.

The common conception that we have to mitigate more at higher levels of atmospheric CO₂ is based on the convexity of damages in global temperature increase. Figure 2 shows that GAUVAL has such a convex damage function, yet, optimal mitigation does not increase in the prevailing CO₂ concentration. The reason lies in the radiative forcing equation (4): the higher the CO₂ concentration, the less does an additional ton of emissions contribute to further forcing and, thus, warming. The precise physics of the underlying logarithm is slightly more complicated, but a simple intuition is as follows. 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. As a consequence, an additional unit of CO₂ emissions has a much lower warming impact than the first unit of anthropogenic emissions. GAUVAL models explicitly the implicit assumptions of Golosov et al. (2014) and Gerlagh & Liski (2012) that the convexity of the damage curve and the concavity of the radiative forcing equation partially offset each other. In contrast to the earlier papers, GAUVAL directly employs the carbon cycle of one of the most wide-spread integrated assessment models, explicitly uses the physical radiative forcing equation, and matches the forcing induced temperature dynamics better than most integrated assessment models. I conclude that the finding might be surprising at first sight, but it is not unreasonable.

In addition, the optimal mitigation policy does not depend on the prevailing temperature level, despite higher temperatures causing higher marginal damages. The reason is that the long-term equilibrium temperature is determined entirely by the GHG concentrations, and a higher temperature level at a given CO₂ concentration implies less warming in the future. GAUVAL shows that this finding prevails in a model that nicely replicates the temperature dynamics of state of the art climate models (Figure 3). 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). GAUVAL states that the marginal social damage curve for CO₂ emissions is flat. In consequence, taxes not only minimize the welfare cost under technological uncertainty and asymmetric information compared to a cap and trade system, but they even eliminate these welfare costs. The marginal damage curve would gain a non-trivial slope if the model was to depart from the assumption of an intertemporal elasticity of substitution of unity. Deterministic estimates usually suggest values smaller than unity. However, the long-run risk literature forcefully argues for an intertemporal elasticity of substitution larger than unity (and disentangled from risk attitude). The logarithmic middle ground stays reasonable. In particular, it is just as easy to argue for a slightly falling marginal damage curve as it is to argue for a slightly increasing marginal damage curve.⁹ The result is convenient for the policy maker: set the

⁹The marginal damage curve would also be negatively sloped if the damage function was less convex,

emission price and let the market take the quantity response. The result is also convenient for the economist: optimal mitigation policy does not require knowledge of the mitigation technology frontier.

3.3 Discounting and Mass Conservation

Optimal economic policy implies that we have to live with the consequences of historic overindulgence in carbon because our mitigation effort is independent of past emissions. What makes it worse: carbon does not decay. Carbon only cycles through the different reservoirs; the fact that some of it eventually turns into limestone is negligible for human planning horizons. A model comparison of scientific carbon cycle models found 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.¹⁰

This implication of mass conservation of carbon has an immediate and important impact on the optimal carbon tax.

Proposition 3 *A carbon cycle (equation 3) satisfying mass conservation of carbon implies a factor $(1 - \beta)^{-1}$, approximately proportional to $\frac{1}{\rho}$, in the closed-form solution of the SCC (equation 11).*

In particular, the SCC approaches infinity as the rate of pure time preference ρ approaches zero.¹¹ I briefly point out how the result changes if I had not normalized population to unity. I assume that the social welfare function is population-weighted per capita consumption and that population grows at the factor $G = \exp(g)$. Then, the root and the factor in equation (11) change to $(1 - \beta G)^{-1} \approx \frac{1}{\rho - g}$. The SCC becomes even more sensitive to the rate of pure time preference. The intuition is that population weighted per-capita consumption puts additional weight on future generations that are more numerous, acting as a reduction of time preference. As is well-known, the value function no longer converges as $\rho \rightarrow g$. Note that, in contrast to the SCC, the temperature's shadow value does not have a root $(1 - \beta)^{-1}$. The matrix σ does not have a unit eigenvalue because the planet exchanges

resembling more closely that of the DICE model. The intuition is that the logarithm in the radiative forcing equation is very strong, and that the underlying saturation in the CO₂ absorption spectrum can outweigh the damage convexity.

¹⁰The 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.

¹¹The present objective function and the dynamic programming equation are not well-defined in the limit of a zero rate of pure time preference. However, the statement holds in that for any $n \in \mathbb{N}$ there exists a strictly positive pure rate of time preference ρ such that $SCC(\rho) > N$.

heat with outer space.¹² The long-run temperature responds to changes of radiative forcing without hysteresis, i.e., without path dependence. Appendix C.1 illustrates Proposition 3 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.

It is well-known that the consumption discount rate plays a crucial role in valuing long-run impacts. The present finding is different. In GAUVAL (as in DICE), the economic impact of climate change grows proportional to output. Given the logarithmic utility specification, future economic growth does not effect the present SCC. The usual consumption discount rate argument based on the Ramsey equation does not apply. Yet, the SCC is extremely sensitive to the rate of pure time preference. It is a widely held belief in the integrated assessment community that it is of little importance how we calibrate the constituents of the consumption discount rate, as long as pure time preference and the growth-based component add up to the same overall consumption discount rate (Nordhaus 2007). The present finding fleshes out the shortcoming of this consumption discount rate based reasoning.

To illustrate the SCC's sensitivity to pure time preference, I reduce the investment-rate-implied annual rate $\rho = 1.75\%$ to a value of $\rho = 1\%$. The SCC increases to 93 USD per ton C or 25.5 USD per ton CO₂. Further reducing the rate of pure time preference to the value of $\rho = 0.1\%$ employed in the Stern (2008) Review results in an optimal carbon tax of 660 USD per ton C and 180 USD per ton CO₂. The Stern Review justified its low pick of the rate of pure time preference by normative reasoning, market failure, and a dual role of individuals who might behave differently on the market compared to large-picture policy decisions (Hepburn 2006).

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. In addition, the present model, like other IAMs, does not explicitly model the actual portfolio of risky investments and, yet, calibrates to overall investment and the Ramsey equation. In an asset pricing context, Bansal et al. (2012) calibrate the pure rate of 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 significantly 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 he discusses some implications for climate change evaluation.

¹²Temperature is an intensive quantity and a conservation of "heat" would imply that the rows of the matrix σ added to unity. However, the atmospheric layer constantly exchanges heat with outer space. Formally, the subtraction of σ^{forc} which implies that the first row of the matrix σ does not add to unity, implying that the largest eigenvalue of the matrix is smaller than unity and historic influences are damped.

Disentangling the different contributions to the SCC for the $\rho = 0.1\%$ case delivers

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

The largest part of the increase, a factor 7, arises from the carbon cycle's contribution. Note that there is no easy way to get the sensitive response to time preference right without separating temperature and carbon dynamics.

4 Carbon Cycle Uncertainty, Learning, and Welfare

This section analyzes the welfare implications of uncertainty governing the carbon flows. I present simple formulas quantifying the welfare impact and the value of uncertainty reduction. First, I model uncertainty as a sequence of persistent shocks drawn from an arbitrary distribution (vector autoregressive model). Second, I present a Bayesian model incorporating epistemological uncertainty and anticipated learning.

4.1 Carbon Sink Uncertainty

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. The limitations in understanding the carbon flows and whether the uptake of the missing sink is permanent or temporary create major uncertainties in predicting future carbon dynamics. The scientific community and governments invest significant sums into the reduction of these uncertainties, including the launching of satellites and new supercomputing facilities. GAUVAL can produce a simple estimate of the welfare costs of these uncertainties and serve as a formal model for quantifying the benefits of uncertainty reduction.

A useful approximation of the carbon cycle uncertainty adds a stochastic flow between the atmosphere and other layers. I modify the carbon cycle's equation of motion (3) to the form

$$\mathbf{M}_{t+1} = \Phi \mathbf{M}_t + (\epsilon_t, -\epsilon_t, 0, \dots, 0)^\top + \mathbf{e}_1 (\sum_{i=1}^I E_{i,t} + E_t^{exo}) , \quad (12)$$

where ϵ_t , $t \in \{0, \dots, \infty\}$ is a sequence of random variables and ϵ_0 has a zero mean. Using DICE 2013's three-layer carbon cycle, the random variable ϵ_t captures the uncertainty in the carbon flow between the atmosphere and the joint upper ocean and biosphere carbon reservoir.

I suggest two conceptually different ways to think about the uncertainty. In the first interpretation, we are worried merely about changes in the carbon flows over time that we cannot predict with certainty. A small persistent shock to ϵ moves the carbon flow and either increases or decreases the sink uptake. 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. In this setting, the present decision maker faces the most uncertainty and we expect to have more knowledge of the carbon flows dynamics in the future, making the system more predictable. I will treat these two cases in the subsequent two sections and compare their economic implications analytically. The final section quantifies the welfare impact and the resulting willingness to pay for a risk reduction.

4.2 Vector Autoregressive (VAR) Uncertainty

I assume a VAR(1) process to model persistent shocks that change the carbon flows in equation (12) in an unpredicted way

$$\epsilon_{t+1} = \gamma\epsilon_t + \chi_t, \quad (13)$$

where $\gamma \leq 1$ and the sequence $\chi_{t,t \in \{0, \dots, \infty\}}$ is independently distributed (and $\epsilon_0 = 0$). Appendix E.1 derives the general welfare difference between the deterministic and the uncertain scenario. Here, I focus on a sequence of identically distributed shocks $\chi_t \sim \chi$ with $\mathbb{E}\chi = 0$. The new shadow value φ_ϵ of the persistent random variable in the carbon cycle is

$$\varphi_\epsilon = \frac{\beta}{1 - \gamma\beta} [\varphi_{M,1} - \varphi_{M,2}]. \quad (14)$$

The welfare cost of a carbon flow change increases in its persistence γ and the shadow value difference between carbon in the atmosphere and in the biosphere-ocean reservoir.

The welfare loss from uncertainty relates closely to the cumulant generating function $G_\chi(z) = \log [\mathbb{E} \exp(z\chi)]$ of the random variable χ . The cumulant generating function is the logarithm of the moment generating function. In contrast to the central moments, cumulants are additive for independent random variables and the i^{th} cumulant is homogenous of degree i under scalar multiplication (cumulant i of $\lambda\chi$ is $\lambda^i \kappa_i$). For most distributions, the cumulant (or the moment) generating functions are tabled and a closed-form solution for the welfare loss follows directly.

Proposition 4 *The vector autoregressive shock process specified in equations (12) and (13) reduces global welfare by*

$$\Delta W^{VAR} = \frac{\beta}{\alpha(1 - \beta)} G_\chi(\alpha\varphi_\epsilon) = \frac{\beta}{\alpha(1 - \beta)} \sum_{i=1}^{\infty} \kappa_i \frac{(\alpha\varphi_\epsilon)^i}{i!}. \quad (15)$$

The first expression states that the welfare loss from carbon flow uncertainty is proportional to the shock's cumulant generating function evaluated at the product of the flow's shadow value φ_ϵ and the intertemporal risk aversion parameter α . The factor $\frac{1}{1-\beta}$ reflects the infinite sum over (stationary) future shocks, and the factor β reflects the one period delay between the shock and its welfare impact.

The second expression for the welfare loss on the r.h.s. expands the function in terms of the random variable's cumulants $\kappa_i, i \in \mathbb{N}$. The first cumulant is the expected value of χ . In the present case where $\kappa_1 = \mathbb{E}\chi = 0$, the second cumulant is the shock's variance and the third cumulant is the third central moment specifying skewness. The higher order cumulants do not coincide exactly with the central moments, but relate closely. The welfare loss is the sum of the stochastic shock's cumulants, each weighted with intertemporal risk aversion and the flow's shadow value taken to the power of the cumulant's order: the expectation is valued independently of risk aversion, the variance proportional to risk aversion, and skewness proportional to risk aversion squared. This basic structure for evaluating the welfare loss of different uncertainties will reappear in all uncertainty specifications.

For a high persistence of the carbon flow shocks (as calibrated in section 4.4), the evaluation of higher moments becomes increasingly sensitive to pure time preference: by equation (14) the shadow value's power φ_ϵ^i is proportional to $(1 - \gamma\beta)^{-i}$. In particular, the welfare loss from a shock's variance and skewness will be more sensitive to the calibration of pure time preference than the welfare loss from the deterministic component (expected flow).

In the case of a mean-zero normally distributed iid shock $\chi \sim N(0, \sigma_\chi^2)$ only the second cumulant $\kappa_2 = \sigma_\chi^2$ differs from zero, and the welfare impact is

$$\Delta W^{VAR,normal} = \frac{\alpha\beta}{1-\beta} \varphi_\epsilon^2 \frac{\sigma_\chi^2}{2} = \frac{\alpha\beta}{1-\beta} \left(\frac{\beta}{1-\gamma\beta} \right)^2 (\varphi_{M,1} - \varphi_{M,2})^2 \frac{\sigma_\chi^2}{2}. \quad (16)$$

The welfare loss is proportional to the shock's variance and the square of the shadow value φ_ϵ and, thus, proportional to the squared difference between the shadow values of carbon in the atmosphere and in the biosphere-ocean reservoir. Its proportionality to $(1 - \gamma\beta)^{-2}$ makes the welfare from an individual shock highly sensitive to high persistence and to low rates of pure time preference.¹³

4.3 Bayesian Uncertainty and Anticipated Learning

This section introduces epistemological uncertainty. It replaces the VAR(1) uncertainty by a Bayesian model, capturing the decision maker's subjective uncertainty as well as her

¹³For an independently distributed shock, where $\gamma = 0$, the welfare loss is not sensitive to time preference if keeping the disentangled Arrow-Pratt risk aversion measure $\frac{\alpha}{1-\beta}$ constant. However, leaving the Arrow-Pratt risk aversion framework, Traeger (2010) argues that α by itself is a measure of intrinsic risk aversion. The binary lotteries in Appendix D illustrate such this measure of risk aversion. Then, the expression in equation (16) is even more sensitivity to time preference.

anticipation of learning over time. Closely related Bayesian learning models have first been used in integrated assessment of climate change by Kelly & Kolstad (1999) in a numeric application to climate sensitivity uncertainty and by Karp & Zhang (2006) in a stylized semi-analytic application to damage uncertainty. Now, a subjective prior governs the uncertain carbon flow

$$\epsilon_t \sim N(\mu_{\epsilon,t}, \sigma_{\epsilon,t}^2), \mu_{\epsilon,0} = 0.$$

The normally distributed prior has an unknown mean and a known variance.

The equations of motion are subject to an objective stochastic shock $\nu_t \sim N(0, \sigma_{\nu,t}^2)$, which can also be interpreted as measurement error. This stochasticity prevents the decision maker from learning the prior's mean from a single observation. Recently, three satellites were launched to reduce the carbon flow measurement errors ν_t (one of whom dropped straight into the Arctic sea). But learning is not limited to future observation. Given the availability of historic data, learning also takes place through the advances in fundamental scientific knowledge and supercomputing. Thus, I interpret $\sigma_{\nu,t}$ merely as a parameter determining the speed of learning. Appendix E.1 spells out the details of the model including the updating equations for the prior.

The new shadow value φ_μ of a shift in the mean carbon flow between atmosphere and ocean is

$$\varphi_\mu = \frac{\beta}{1-\beta} [\varphi_{M,1} - \varphi_{M,2}]. \quad (17)$$

This shadow value coincides with the shadow value of a *perfectly persistent* shock (equation 14).

Proposition 5 *Bayesian epistemological carbon flow uncertainty with anticipated learning results in the welfare loss*

$$\Delta W^{Bayes} = \sum_{t=0}^{\infty} \beta^{t+1} \frac{\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2}{2} \alpha (\varphi_{M,1} - \varphi_{M,2})^2 \left(\frac{1}{1-\beta} \right)^2 \Omega_t^2 \quad (18)$$

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

which assumes a normal-normal conjugate prior model governing the learning about the uncertainty ϵ_t in equation (12).

As in the case of a normally distributed autoregressive shock, the welfare loss is proportional to (intertemporal) risk aversion and the squared difference between the shadow values of carbon in the atmosphere and in the biosphere-ocean sink. The relevant variance in the case of learning is the prior's variance plus the variance of the periodic shock $\sigma_{\epsilon,t}^2 + \sigma_{\nu,t}^2$. The prior's variance is declining as the decision maker improves her estimate of the true carbon flows. As a result of this non-stationarity, an infinite discounted sum replaces the factor $\frac{1}{1-\beta}$ observed in the VAR model of equation (16).

The main difference between the Bayesian learning model and the VAR model is the role of the uncertainty's persistence and the resulting sensitivity to time preference. *Initially, learning updates* in the Bayesian framework act as *perfectly persistent shocks*. This effective persistence makes the Bayesian decision maker much more sensitive to pure time preference with respect to uncertainty and learning shocks in the first few periods. Intuitively, while the prior is not yet sharply focused on any given value, the Bayesian decision maker updates her long-run belief over the expected carbon flows. If she is patient, this long-run update moves her welfare significantly. Formally, this finding is reflected in the factor $\left(\frac{1}{1-\beta}\right)^2$ replacing the factor $\left(\frac{1}{1-\gamma\beta}\right)^2$ of the VAR shock model (perfect persistence is $\gamma = 1$).

Over time, the Bayesian decision maker learns the carbon flow dynamics and eliminates epistemological uncertainty. The remaining uncertainty is merely noise. The shock's no longer carry informational value and their welfare impact corresponds merely to that of independently distributed stochastic shocks with no persistence. Therefore, later periods contribute little to the welfare loss from uncertainty and are hardly sensitive to pure time preference. Formally, this finding is reflected in the factor Ω_t . It is a weighted mean of unity and $1 - \beta$. Initially, the variance of the prior $\sigma_{\epsilon,t}^2$ dominates, and $\Omega \approx 1$. Over time, the decision maker learns the subjectively uncertain part of the carbon flow uncertainty. As $\sigma_{\epsilon,t}^2$ falls, the weight on $1 - \beta$ increases and eventually $\Omega_t^2 \rightarrow (1 - \beta)^2$, offsetting the term $\left(\frac{1}{1-\beta}\right)^2$. Finally, the Bayesian decision maker is uncertain about the carbon flow already in the present period. Therefore, the formula summarizing the welfare loss in Proposition 5 misses one power of the discount factor β compared to the VAR model of equation (16).

4.4 Quantifying the Welfare Impact of Carbon Cycle Uncertainty

The quantitative assessment of carbon cycle uncertainty is based on the DICE 2013 business as usual scenario and Joos et al. (2013), who subject 18 different carbon cycle models to a 100Gt and a 5000Gt carbon pulse and track their responses for up to 3000 years. The larger shock corresponds to 500 years of present day emissions and results in a medium run (100-1000 years) cross-model standard deviation of atmospheric carbon of approximately 500Gt. Joos et al.'s (2013) comparison study is subject to common bias, and this standard deviation is best interpreted as a lower bound of actual uncertainty. In the VAR model, Joos et al.'s (2013) simulations suggest a persistence in equation (16) of $\gamma = 0.997$, i.e., highly intertemporally correlated uncertainty governing the flow of carbon between the atmosphere and the sinks. I evaluate all carbon cycle uncertainty scenarios with a risk aversion coefficient of $\alpha = -1.5$ (see Appendix D).

For the auto-regressive shock scenario of section 4.2, I assume a standard deviation of the decadal shock of $\sigma_\chi = 20$ Gt per decade, which builds up into the 500 Gt forecast standard deviation suggested by Joos et al. (2013) approximately by the year 2300. Figure 4, left panel, evaluates the resulting carbon concentration along the business as usual scenario of DICE 2013. Appendix E.3 discusses variations. This uncertainty translates into a present

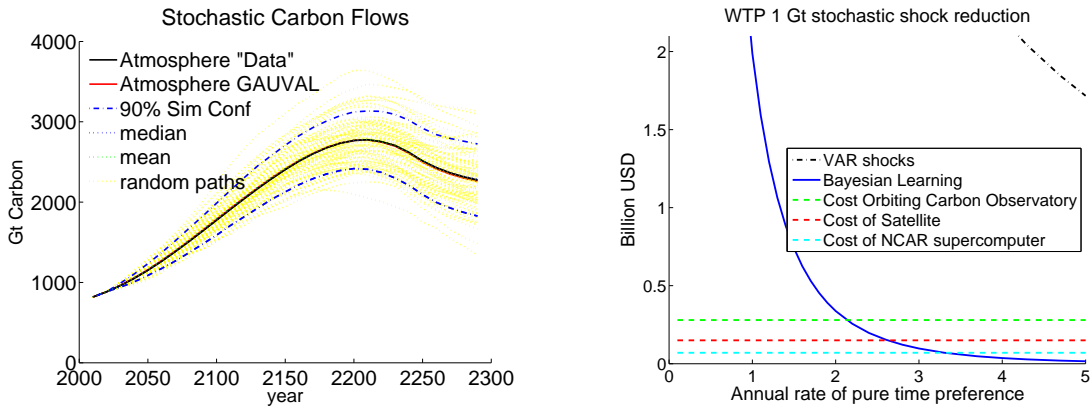


Figure 4: The graph on the left shows the evolution of atmospheric carbon for the DICE 2013 business as usual emission scenario. Decadal shocks with standard deviation of $\sigma_\chi = 20$ Gt per decade change the flow between the atmosphere and the carbon sinks with a persistence of $\gamma_M = 0.997$ calibrated to carbon cycle comparison study by Joos et al.’s (2013). 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. 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).

value welfare loss of 110 billion USD. For a comparison of magnitude, the annual NASA budget is about 20 billion USD.

For the case of epistemological uncertainty, I assume that the Bayesian prior has a standard deviation of $\sigma_{\epsilon,0} = 20$ Gt per decade, and that the “measurement error” is $\sigma_\nu = 10$ 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 resulting welfare loss from carbon cycle uncertainty in this Bayesian learning scenario is approximately 30 billion USD. This value is most sensitive to the initial prior. Lowering initial uncertainty to $\sigma_{\epsilon,0} = 10$ 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.

In the case of an annual rate of pure time preference of $\rho = 0.1$ the uncertainty contributions increase to 2 trillion USD in the case of VAR shocks and to 60 trillion USD, or 73% of world output, in the case of Bayesian learning. Now, the Bayesian learning scenario implies the (much) higher welfare loss. It is a manifestation of the higher sensitivity to pure time preference that I derived in the previous section: the information gain from early shocks has an impact comparable to a perfectly persistent autoregressive shock, and the welfare loss from such shocks increases strongly in the decision maker’s patience.

It is fair to 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 and the remaining welfare costs are born directly from risk aversion. 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). For the patient decision maker, the welfare loss from uncertainty carefully catches up to a similar order of magnitude.¹⁴

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 ($\sigma_\chi = 10\text{Gt} \rightarrow 9\text{Gt}$).¹⁵ 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 (American) 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 GAUVAL.¹⁶ In contrast, 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 110 billion USD, emphasizing again the particularly high sensitivity of the Bayesian uncertainty model to pure time preference.

5 Temperature Uncertainty

The present section discusses the basics of temperature uncertainty. I explain why the normal-normal learning model would not be reasonable in the case of temperature uncertainty and construct an alternative model combining epistemological uncertainty with persistent shocks. I quantify a lower bound for the welfare loss from uncertainty over the temperature response to a given CO₂ concentration (climate sensitivity).

¹⁴The 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.

¹⁵Reducing 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 much more valuable than a reduction of measurement error that accelerates learning.

¹⁶NASA'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.

5.1 On Temperature Tails

So far, I have assumed that a doubling of the CO₂ concentrations from its pre-industrial concentration of 280 ppm to 560 ppm yields a medium-term temperature increase of 3C. 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 lower than the corresponding equilibrium increase because of the atmosphere-ocean temperature interaction discussed in section 2.2. The implied 3C warming in the (medium-run) equilibrium is little more than a best guess. The value depends on a set of uncertain feedback processes that either increase or decrease the initial warming. For example, 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).

Meinshausen et al. (2009) collect 20 estimates of the probability distributions governing the temperature increase from a doubling of the CO₂ concentration, the so-called climate sensitivity. These estimates derive from different research groups and use a variety of methodological approaches. My evaluation of temperature uncertainty relies on the average distribution assigning equal weight to each approach. The climate sensitivity distribution governing global warming is positively skewed, i.e., it exhibits a pronounced right tail. I refer to Appendix F.2 and Figure 8 for details. The support of these distributions spans from no warming to a 10C warming. The climate sensitivity distribution governing global warming is positively skewed, i.e., it exhibits a pronounced right tail. A warming above 10C is possible but lacks probabilistic estimates.

GAUVAL'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 section 4.2, 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 (2009) "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 (2009) can blow up the welfare loss through its translation into economic damages.¹⁷ In a DICE-style stochastic IAM Kelly & Tan (2013) show that fat tails on climate sensitivity are accompanied by relatively quick learning. Such quick learning has, first, not been observed and, second, would make the tails of climate sensitivity less important. These results lessen the concern about fat-tailed uncertainty over climate sensitivity. However, in contrast to Weitzman's result, the present reasoning does not

¹⁷Figure 2 shows that for temperature increases up to 12C, GAUVAL'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.

rely on the fatness of the climate sensitivity tail. Thin-tailed uncertainty interacts with temperature dynamics and the convexity of damages and continues to emphasize the importance of using stochastic rather than deterministic models for climate change evaluation.

These results rely on the extrapolation of different functional forms that capture our best guess of what happens at moderate degrees of warming. However, no economic model to date has incorporated a reasonable quantitative characterization of what happens to life and welfare on planet Earth for a 20C or 30C global temperature increase. Should we conclude that we commit all of our economic resources to fighting global warming? No.¹⁸ It is already daring to evaluate damages from a warming of up to 10C with state of the art integrated assessment models. To evaluate even higher temperature and damage scenarios, we should *not* rely on the extrapolation of functional forms, but build a separate model aiming directly at the quantification of scenarios that lie far from experience and imagination. What GAUVAL can do is estimate a lower bound of the welfare loss from uncertainty given a warming range for which damages and temperature dynamics seem somewhat reasonable.

5.2 A VAR Model with Epistemological Uncertainty

Temperature uncertainty has a major epistemological component. The only Bayesian (conjugate prior) model compatible with the risk aversion enhanced linear-in-state model seems to be the normal-normal learning model that I employed in section 4.3. For temperature uncertainty, I cannot reasonably assume a normal distribution.¹⁹

I introduce a combined epistemological and vector autoregressive shock model of uncertainty that disposes of the straight-jacket imposed by a normal distribution. This step comes at the expense of no longer being a strictly Bayesian model. I characterize an arbitrary probability distribution through its (countably infinite) sequence of cumulants, instead of relying on a parametric (or even conjugate) class of probability distributions. This approach further illustrates the usefulness of the cumulant-based approach, and it allows me to employ the numeric distribution derived in the scientific literature for a quantitative estimate of the welfare impact.

I represent uncertainty governing the temperature's equation of motion as a random contribution ϵ_t^τ to incremental warming, changing equation (8) to the form

$$\tau_{t+1} = \sigma \tau_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} e_1 + \epsilon_t^\tau e_1 . \quad (19)$$

¹⁸Also technically, the result hinges on a crucial assumption that makes it inadequate for the evaluation of high temperature tails. Representations of rational preferences by expected utility, including the present representation, require bounded utility functions or alternative assumptions with similar consequences (von Neumann & Morgenstern 1944, Kreps 1988). The “dismal result” depends crucially on *not* meeting this rationality assumption when the damage distribution approaches full output loss with high density.

¹⁹A normal on temperature in degree Celsius would ignore positive skewness and, yet, results would be driven by tails far from what I can reasonably calibrate. A normal on transformed temperature τ_1 would assign significant weight to large negative values of the atmospheric temperature increase T_1 in degree C.

The random variable ϵ_t^τ captures epistemological uncertainty in period t . In analogy to the Bayesian model, only the distribution (not the realization) of ϵ_t^τ is known to the decision maker in period t . Future shocks to this distribution capture both actual shocks as well as informational updates (see also footnote 20 below). I characterize the distribution of ϵ_t^τ by its cumulants $\kappa_{i,t}, i \in \mathbb{N}$. Initial epistemological uncertainty is given by $\kappa_{i,0}, i \in \mathbb{N}$, and the cumulants follow the equations of motion

$$\kappa_{i,t+1} = \gamma^i \kappa_{i,t} + \chi_{i,t}^\tau, \quad (20)$$

$0 \leq \gamma \leq 1$, for all $i \in \mathbb{N}$. The special case where $\kappa_i = 0$ for $i > 2$ relates closely to the Bayesian learning model in section 4.3.²⁰ In the absence of shocks ($\chi_{i,t}^\tau = 0$ for all i), epistemological uncertainty decays at rate γ : ϵ_{t+1} is then distributed as $\gamma \epsilon_t$ (the i^{th} cumulant is homothetic of degree i). I assume that the stochastic shocks $\chi_{i,t}^\tau$ are independently and identically distributed (see Appendix F.1 for the non-stationary case). These shocks modify epistemological knowledge and uncertainty. In the normally distributed Bayesian learning model, the uncertainty over the mean falls monotonically and quickly. In the VAR model, the persistent shocks remain constant over time and keep building up into forecast uncertainty. The present model combines these features. In addition, the model in equation (20) permits shocks to the variance $\chi_{2,t}^\tau$ (stochastic volatility), making the speed of learning itself uncertain (a feature also present in “non-normal” Bayesian learning models).²¹

Proposition 6 *The combined VAR-epistemological model for temperature uncertainty results in the welfare loss*

$$\Delta W^{\text{temp}} = \underbrace{\sum_{i=1}^{\infty} \varphi_{\kappa,i} \kappa_{i,0}}_{\text{Epistemological}} + \underbrace{\frac{\beta}{\alpha(1-\beta)} \sum_{i=1}^{\infty} G_{\chi_i^\tau}(\alpha \varphi_{\kappa,i})}_{\text{Future Shocks}} \quad (21)$$

where the cumulants $\kappa_{i,0}$ characterize the present epistemological uncertainty distribution (governing ϵ_0^τ in equation 19) and have the shadow values

$$\varphi_{\kappa,i} = \frac{\beta}{1 - \beta \gamma^i} \frac{(\alpha \varphi_{\tau,1})^i}{i! \alpha}, \quad (22)$$

and the random variables χ_i^τ characterize iid shocks to the i^{th} cumulant of this epistemological distribution (equation 20).

²⁰The Bayesian model would exhibit a time dependent parameter $\gamma_{i,t}$ in the equation of motion $\kappa_{i,t+1} = \gamma_{i,t} \kappa_{i,t} + \chi_{i,t}^\tau$ and require an additional non-persistent iid shock in equation (19) (which has little impact on the welfare evaluation). The parameter $\gamma_{2,t} = \sigma_{\epsilon,t}^2$ can be chosen to reproduce the learning of the Bayesian model (plus stochasticity ν_t), and the shocks $\chi_{1,t}^\tau$ would have to match the Bayesian updates of the mean.

²¹Not all combinations of shocks to the cumulants will lead to well-defined probability distributions of ϵ_t . However, in such a case, the cumulants can still be considered a reasonable approximation of a closely related distribution.

The first contribution to the welfare loss in equation (21) derives from the epistemological uncertainty in the present ($t = 0$). The shadow values directly capture the corresponding welfare loss that results from a unit increase of a given cumulant. The shadow value $\varphi_{\kappa,1} = \frac{\beta}{1-\beta\gamma}\varphi_{\tau,1}$ of a mean shift in (transformed) temperature flow is equivalent to the shadow value φ_ϵ in the VAR model (equation 14, with $\varphi_{\tau,1}$ replacing $\varphi_{M,1}-\varphi_{M,2}$). It states the welfare loss incurred from a shift of the mean temperature flow. The shadow value $\varphi_{\kappa,2} = \frac{\beta}{1-\beta\gamma^2}\frac{\alpha}{2}\varphi_{\tau,1}^2$ of a unit increase of the variance combines characteristics of both the normal VAR shock welfare loss in equation (15) and the Bayesian model's welfare loss in equation (18). The welfare impact prevails immediately and its impact ceases over time as in the Bayesian model, and the decay is parameterized by the simple persistence parameter γ resembling (though not coinciding with) that of the VAR model.

Higher order cumulants contribute to the welfare loss proportional to the power $i-1$ of intertemporal risk aversion and the power i of the generalized temperature's shadow value, a finding also observed in section 4.2's VAR setting. However, in contrast to the stationary VAR shocks in equation (15), the higher order moments of the epistemological distribution have a welfare impact that is less sensitive to pure time preference compared to the welfare impact of the distribution's variance: epistemological uncertainty decays faster for higher order cumulants thanks to the cumulants' homogeneity of degree i .

The second contribution to the welfare loss in equation (21) derives from future shocks. These future shocks capture both autoregressive shocks as well as updates to the epistemological distribution. Due to the simplified setup, the persistence of an autoregressive shock coincides with the persistence of epistemological uncertainty γ . The first summand of the r.h.s. sum in equation (21) captures shocks to the mean. The cumulant generating function $G_{\chi_1^\tau}$ of the iid shock to the mean governs the welfare impact. It is evaluated at the (intertemporal) risk aversion weighted shadow value of a unit shift of the mean. This welfare loss is equivalent to the one observed in equation (15) for VAR shocks to the carbon flow. In particular, the i^{th} cumulant of the shock χ_1^τ contributes proportional to the power $i-1$ in risk aversion and to the power i in the temperature's shadow value. Note that the cumulants of the shock χ_1^τ differ from the cumulants $\kappa_{i,\tau}$, $i \in \mathbb{N}$, that capture the uncertainty ϵ_t^τ prevailing in period t .

The second summand of the r.h.s. sum in equation (21) captures shocks to the variance of the distribution ϵ_τ . This type of uncertainty was absent in the pure VAR shock model and the normal-normal Bayesian learning model. It is similar to a stochastic volatility model. Here, the future uncertainty's variance is itself subject to random shocks captured by the random variable χ_2^τ . The normal-normal Bayesian learning model is special in that the variance of the posterior falls deterministically. More general Bayesian learning models share the present feature that some realizations of shocks can increase future uncertainty. The simplest shock is a mean-zero normal shock to the variance, which is fully characterized by χ_2^τ 's variance. This shock causes a welfare loss proportional to the third power of risk aversion and the

forth power of the temperature's shadow value. More generally, an i^{th} order shock to the j^{th} cumulant of the uncertainty distribution causes a welfare loss that is proportional to risk aversion in the power $i \cdot j - 1$ and to the temperature's shadow value to the power $i \cdot j$.

Finally, the sensitivity to pure time preference increases in the order of the shock's moment but decreases in the order of the uncertainty distribution's moment that is being shocked. The reduction in sensitivity to pure time preference for higher order moments is analogous to the case of the epistemological contribution discussed above and it vanishes for perfect persistence $\gamma \rightarrow 1$. The increasing sensitivity to pure time preference in the order of the shock's moment is analogous to the case of mere VAR shocks discussed in section 4.2.

5.3 Quantifying the Welfare Impact of Temperature Uncertainty

Once more, I employ DICE's business as usual (BAU) emission scenario for a quantitative assessment. The solid (deterministic) line in Figure 4 depicts the resulting carbon concentration. I employ the average of Meinshausen et al.'s (2009) survey of climate sensitivity distributions, whose mean I shift to the 3C best guess, focusing on the impact of uncertainty only. Greenhouse gas concentrations significantly exceed a doubling of CO_2 along the DICE BAU scenario, and I scale the temperature shocks to the corresponding forcing. Initial epistemological uncertainty ϵ_0^τ and the distribution of the shocks to the mean $\chi_{1,t}$ are both chosen to reproduce climate sensitivity uncertainty if concentrations are double the pre-industrial level. Epistemological uncertainty ϵ_0^τ prevails immediately, whereas the shock uncertainty builds up over time.

I make the following four assumptions in the spirit of finding a lower bound of the welfare loss. First, I calibrate the shock uncertainty to produce the climate sensitivity distribution in the infinitely long run. Second, I split the overall climate sensitivity 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 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 to CO_2 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

GAUVAL is an integrated assessment model of climate change that matches scientific climate models just as well as numeric counterparts used in policy advising. It derives the optimal carbon tax and welfare loss in closed form. GAUVAL 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 allows for a better 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 CO_2), 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%). Like Golosov et al.’s (2014) model, GAUVAL 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. 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 compared to the 1.75% calibrated here following a standard approach. These more sophisticated approaches support rates as low as the 0.1%, which the Stern 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 GAUVAL to advance our understanding of the welfare implications of uncertainty in climate change, distinguishing between epistemological uncertainty and shocks to the climate system. 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 exactly 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. An accompanying paper analyzes the response of the optimal carbon tax to climate and economic uncertainties. GAUVAL 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. GAUVAL also lends itself to a normatively motivated analysis. GAUVAL'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* .
- 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.
- 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.
- Crost, B. & Traeger, C. P. (2014), 'Optimal co2 mitigation under damage risk valuation', , *Nature Climate Change* **4**, 631–636.

- 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.
- 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.
- Hepburn, C. (2006), Discounting climate change damages: Working notes for the Stern review, Working note.
- Hoel, M. & Karp, L. (2002), ‘Taxes versus quotas for a stock pollutant’, *Resource and Energy Economics* **24**, 367–384.
- 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. P. (2014), ‘Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings’, *European Economic Review* **69**(C), 104–125.
- 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. & Zhang, J. (2006), ‘Regulation with anticipated learning about environmental damages’, *Journal of Environmental Economics and Management* **51**, 259–279.

- Kelly, D. L. & Kolstad, C. D. (1999), ‘Bayesian learning, growth, and pollution’, *Journal of Economic Dynamics and Control* **23**, 491–518.
- Kelly, D. L. & Tan, Z. (2013), ‘Learning and climate feedbacks: Optimal climate insurance and fat tails’, *University of Miami Working paper*.
- Kreps, D. (1988), *Notes on the Theory of Choice*, Westview Press, Boulder.
- 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.
- Moss, R., Babiker, M., Brinkman, S., Calvo, E., Carter, T., Edmonds, J., Elgizouli, I., Emori, S., Erda, L., Hibbard, K., Jones, R., Kainuma, M., Kelleher, J., Lamarque, J. F., Manning, M., Matthews, B., Meehl, J., Meyer, L., Mitchell, J., Nakicenovic, N., O’Neill, B., Pichs, R., Riahi, K., Rose, S., Runci, P., Stouffer, R., van Vuuren, D., Weyant, J., Wilbanks, T., van Ypersele, J. P. & Zurek, M. (2007), *Towards New Scenarios for Analysis of Emissions, Climate Change, Impacts, and Response Strategies.*, Intergovernmental Panel on Climate Change, Geneva.
- 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.

- Rabin, M. (2000), 'Risk aversion and expected-utility theory: A calibration theorem', *Econometrica* **68**(5), 1281–1292.
- Schneider, M., Traeger, C. P. & Winkler, R. (2013), 'Trading off generations: Infinitely lived agent versus olig', *European Economic Review* **56**, 1621–1644.
- Stern, N. (2008), 'The Economics of Climate Change', *American Economic Review* **98**(2), 1 – 37.
- Traeger, C. (2010), 'Intertemporal risk aversion – or – wouldn't it be nice to know whether Robinson is risk averse?'. CUDARE Working Paper No. 1102.
- Traeger, C. (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-sided 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* .
- 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), 'Unexpected 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. (2009), 'On modeling and interpreting the economics of catastrophic climate change', *The Review of Economics and Statistics* **91**(1), 1–19. 06.
- Weitzman, M. L. (2010), 'Ghg targets as insurance against catastrophic climate damages', *NBER Working Paper* (16136).

Appendix

Part I - The Deterministic Model

A General Capital Depreciation

Equation (2) 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 (2) 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

$$K_{t+1} = Y_t[1 - D_t(T_{1,t})](1 - x_t) \left[\frac{1 + g_{k,t}}{\delta_k + g_{k,t}} \right] . \quad (23)$$

For full depreciation $\delta_k = 1$ the squared bracket is unity and equation (23) coincides with equation (2) 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 (2) with equation (23) 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 (2) and (23) 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 GAUVAL’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*).²² It also abandons the assumption of logarithmic utility, further stacking the cards against GAUVAL’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 2.5).

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

B Solution of the Linear-in-State Model

B.1 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 $\tau_{i,t} \equiv \exp(\xi_i T_{i,t})$,

²²The 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.

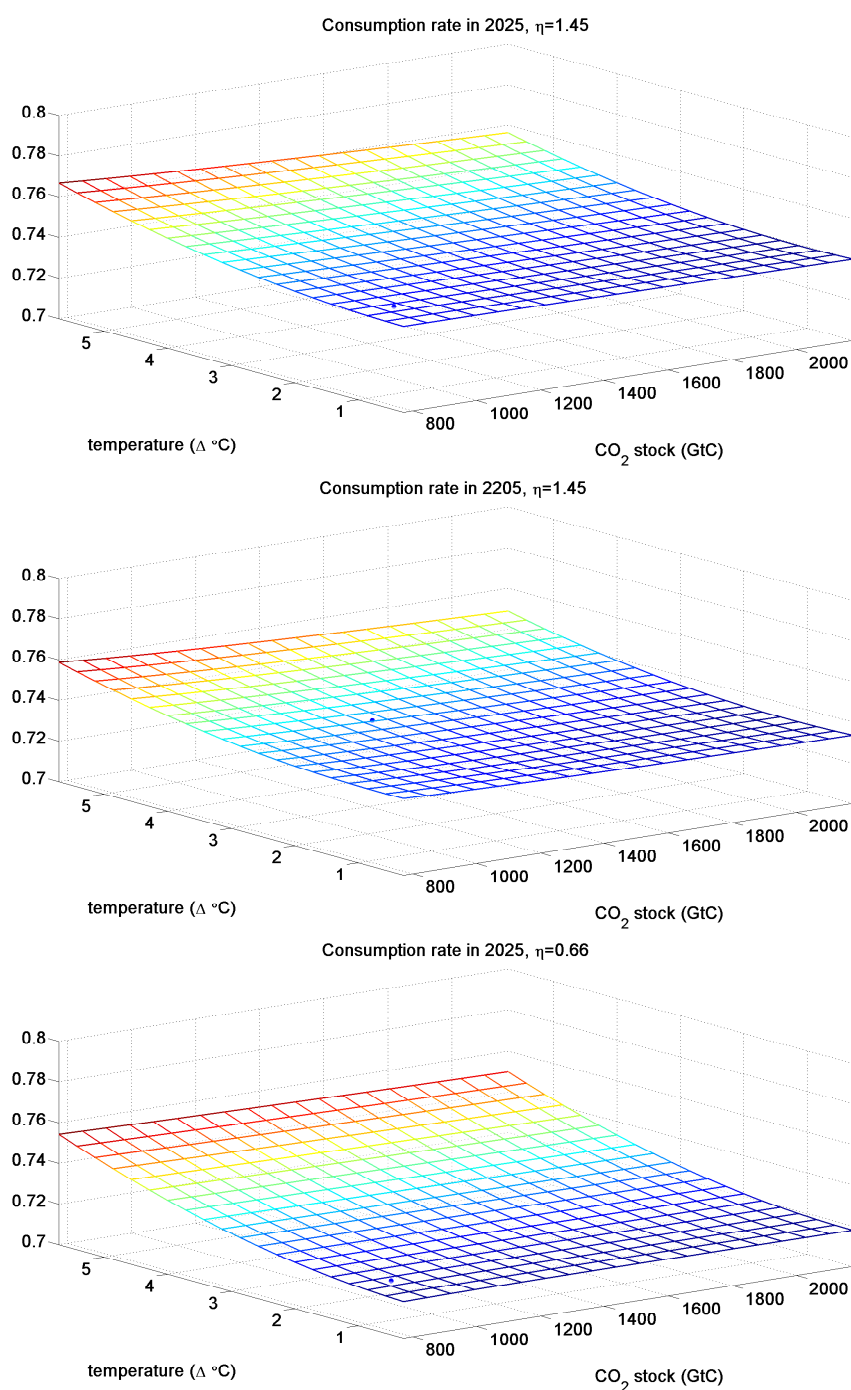


Figure 5: The graphs analyze the climate (in-)dependence of the optimal consumption rate x^* in the wide-spread DICE model, relying on the control rules of the recursive implementation by Traeger (2012b). The first two graphs assume the DICE-2013 value $\eta = 1.45$, the third graph follows the long-run risk literature with $\eta = \frac{2}{3}$. The blue dot in each graph indicates the expected optimal control and prevailing temperature- CO_2 combination along the optimal policy path in the given year.

$i \in \{1, \dots, L\}$. 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

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

The Bellman equation in terms of the transformed state variables is

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \max_{x_t, \mathbf{N}_t} \log x_t + a_t + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} \\ + \nu \log E_t - \xi_0 \tau_t + \xi_0 + \beta V(k_{t+1}, \boldsymbol{\tau}_{t+1}, \mathbf{M}_{t+1}, \mathbf{R}_{t+1}, t+1), \quad (24)$$

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

$$k_{t+1} = a_t + \kappa k_t + (1 - \kappa - \nu) \log N_{0,t} + \nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t)), \\ - \xi_0 \tau_{1,t} + \xi_0 + \log(1 - x_t) \quad (25)$$

$$\mathbf{M}_{t+1} = \boldsymbol{\Phi} \mathbf{M}_t + \left(\sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} \right) \mathbf{e}_1. \quad (26)$$

I transform the temperature's equation of motion (5) 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

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

still using $\sigma_{i,l+1} = 0$ 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

$$\boldsymbol{\tau}_{t+1} = \boldsymbol{\sigma} \boldsymbol{\tau}_t + \sigma^{forc} \frac{M_{1,t} + G_t}{M_{pre}} \mathbf{e}_1 . \quad (27)$$

Finally, the equation of motion for the resource stock is

$$\mathbf{R}_{t+1} = \mathbf{R}_t - \mathbf{E}_t^d , \quad (28)$$

The underlying constraints are

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

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

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.

B.2 Proof of Proposition 1

I start by showing that the affine value function

$$V(k_t, \boldsymbol{\tau}_t, \mathbf{M}_t, \mathbf{R}_t, t) = \varphi_k k_t + \boldsymbol{\varphi}_M^\top \mathbf{M}_t + \boldsymbol{\varphi}_\tau^\top \boldsymbol{\tau}_t + \boldsymbol{\varphi}_{R,t}^\top \mathbf{R}_t + \varphi_t \quad (29)$$

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 (25)-(28) 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 in equation (29) and the next period

states (equations 25-28) into the (deterministic) Bellmann equation (24) delivers

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

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

$$\frac{1}{x} - \beta \varphi_k \frac{1}{1-x} = 0 \quad \Rightarrow \quad x^* = \frac{1}{1 + \beta \varphi_k} . \quad (30)$$

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 $\mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \varphi_M, \varphi_{R,t+1})$. 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

$$\begin{aligned}
& (\varphi_k - (1 + \beta\varphi_k)\kappa)k_t + \left(\varphi_M^\top - \beta\varphi_M^\top \Phi - \beta\varphi_{\tau,1} \frac{\sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top \right) \mathbf{M}_t \\
& + (\varphi_\tau^\top - \beta\varphi_\tau^\top \sigma + (1 + \beta\varphi_k)\xi_0 \mathbf{e}_1^\top) \boldsymbol{\tau}_t + (\varphi_{R,t}^\top - \beta\varphi_{R,t+1}^\top) \mathbf{R}_t \\
& + \varphi_t = \log x_t^*(\varphi_k) + \beta\varphi_k \log(1 - x_t^*(\varphi_k)) + (1 + \beta\varphi_k)\xi_0 + (1 + \beta\varphi_k)a_t \\
& \quad + (1 + \beta\varphi_k)(1 - \kappa - \nu) \log \mathbf{N}_{0,t}^*(\mathbf{A}_t, \varphi_k, \varphi_M, \varphi_{R,t+1}) \\
& \quad + (1 + \beta\varphi_k)\nu \log g(\mathbf{E}_t(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \varphi_M, \varphi_{R,t+1}))) \\
& \quad + \beta\varphi_{M,1} \sum_{i=1}^{I^d} E_{i,t}(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \varphi_M, \varphi_{R,t+1})) + E_t^{exo} \\
& \quad - \beta\varphi_{R,t+1}^\top \mathbf{E}_t^d(\mathbf{A}_t, \mathbf{N}_t^*(\mathbf{A}_t, \varphi_k, \varphi_M, \varphi_{R,t+1})) \\
& \quad + \beta\varphi_{t+1} .
\end{aligned} \tag{31}$$

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

$$\varphi_k - (1 + \beta\varphi_k)\kappa = 0 \quad \Rightarrow \quad \varphi_k = \frac{\kappa}{1 - \beta\kappa} \tag{32}$$

$$\varphi_M^\top - \beta\varphi_M^\top \Phi - \beta\varphi_{\tau,1} \frac{\sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top = 0 \quad \Rightarrow \quad \varphi_M^\top = \frac{\beta\varphi_{\tau,1}\sigma^{forc}}{M_{pre}} \mathbf{e}_1^\top (\mathbf{1} - \beta\Phi)^{-1} \tag{33}$$

$$\varphi_\tau^\top + (1 + \beta\varphi_k)\xi_0 \mathbf{e}_1^\top - \beta\varphi_\tau^\top \sigma = 0 \quad \Rightarrow \quad \varphi_\tau = -\xi_0(1 + \beta\varphi_k)\mathbf{e}_1^\top (\mathbf{1} - \beta\sigma)^{-1} \tag{34}$$

$$\varphi_{R,t}^\top - \beta\varphi_{R,t+1}^\top = 0 \quad \Rightarrow \quad \varphi_{R,t} = \beta^t \varphi_{R,0} .$$

The initial values $\varphi_{R,0}^\top$ 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 (30) results in the optimal investment rate $x = 1 - \beta\kappa$. From line (31) 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.

B.3 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 (32) for the shadow value of log-capital and (34) for the shadow value of atmospheric temperature (first entry of the vector) into equation (33) characterizing the shadow value of carbon in the different reservoirs delivers

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

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

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

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

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

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

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

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

$$\Delta W^{seq} = \varphi_{M,i} - \varphi_{M,1} = \frac{\beta\varphi_{\tau,1}\sigma^{forc}}{M_{pre}} \left([(\mathbf{1} - \beta\Phi)^{-1}]_{1,i} - [(\mathbf{1} - \beta\Phi)^{-1}]_{1,1} \right). \quad (35)$$

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:

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

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.²³ Appendix C.3 illustrates equation (35) for a two layer carbon cycle and discusses more generally the relation between carbon prices in different reservoirs.

B.4 Proof of Proposition 3

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 $\mathbf{1} - \beta\Phi$ has the root $1 - \beta$. It follows from Cramer's rule (or as an application of the Cayley-Hamilton theorem) that the entries of the matrix $(\mathbf{1} - \beta\Phi)^{-1}$ are proportional to $(1 - \beta)^{-1}$.

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

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

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 3's root $(1 - \beta)^{-1}$ that makes the expression so sensitive to a low rate of pure time preference.

²³Note that the present model does not explicitly model damages from ocean acidification, which would be an interesting and feasible extension.

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

$$\varphi_{M,1} = \beta \varphi_{\tau,1} \sigma^{forc} M_{pre}^{-1} (1 - \beta(1 - \delta_{decay}))^{-1}, \quad (36)$$

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 11) takes the form

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

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 β accompanies each of these transition coefficients as each of them causes a one period delay. Taking the limit of $\beta \rightarrow 1$ confirms that (an analogue to) Proposition 3 does not hold for the temperature system

$$\lim_{\beta \rightarrow 1} \varphi_{\tau,1} = -\xi_0(1 + \varphi_k) [(1 - \sigma)^{-1}]_{11} = -\frac{\xi_0(1 + \varphi_k)}{\sigma_1^{up}} \neq \infty. \quad (37)$$

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.²⁴

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 σ do not sum to unity. As a consequence of the mean structure in equation (5), 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

²⁴I note that the carbon cycle lacks the reduction in persistence deriving from the forcing weight σ_1^{up} . With this observation, equation (37) 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 $\beta \rightarrow 1$) 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.

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

$$\varphi_{M,i} = \beta \frac{\sum_{j \neq i} \varphi_{M,j} \Phi_{j,i} + 1_{i,1} \frac{\varphi_{\tau,1} \sigma_1^{up}}{M_{pre}}}{1 - \beta \Phi_{i,i}}. \quad (38)$$

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 (38) resembles the carbon pricing formula for the carbon decay model discussed in equation (36) 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 (35) for the case of the two layer carbon cycle introduced in section C.1

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

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 - Uncertainty

D 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 (10) 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

$$V_t^* = \exp \left((1 - \beta) \log c_t + \beta \log \left[\mathbb{E}_t V_{t+1}^* \right]^{\frac{1}{\alpha^*}} \right). \quad (39)$$

Defining $V_t = \frac{\log V_t^*}{1 - \beta}$ and rearranging equation (40) delivers

$$V_t = \log c_t + \frac{\beta}{1 - \beta} \log \left[\mathbb{E}_t \exp \left((1 - \beta) V_{t+1} \right)^{\alpha^*} \right]^{\frac{1}{\alpha^*}}. \quad (40)$$

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 (10) stated in the text.

Figure 6 illustrates the quantitative implications of a choice of risk aversion $\text{RRA} = 1 - \alpha$ in the model.²⁵ 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 (10) and into $\text{RRA} \in \{2, 2.5\}$ in Figure 6.

²⁵I 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.

E Carbon Cycle Uncertainty

E.1 Proof of Proposition 4

In the case of persistent carbon sink shocks, the adjustments in the equations of motion (12) and (13) modify or add the following terms to the Bellman equation (10)

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

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 (14) 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

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

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

$$\Delta W^{VAR} = \frac{1}{\alpha} \sum_{t=0}^{\infty} \beta^{t+1} \log \left[\mathbb{E} \exp [\alpha \varphi_\epsilon \chi_t] \right]. \quad (41)$$

For a sequence of identically distributed shocks χ_t , I obtain the welfare cost of uncertainty stated in (15) by evaluating the implied geometric sum in equation (41).

E.2 Proof of Proposition 5

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

$$\begin{aligned} M_{1,t+1} &= (\Phi M_t)_1 + \sum_{i=1}^{I^d} E_{i,t} + E_t^{exo} + \epsilon_t + \nu_t, \\ M_{2,t+1} &= (\Phi M_t)_2 - \epsilon_t - \nu_t. \end{aligned} \quad (42)$$

I model the learning process based on atmospheric carbon observation.²⁶ Rearranging equation (42), the decision maker derives information on ϵ_t from the realizations

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

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

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

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.

These adjustments in the equations of motion imply modifications of the Bellman equation (10) captured by the terms

$$\begin{aligned} \varphi_\mu \mu_{\epsilon,t} + \varphi_t + \dots &= \dots + \beta \varphi_{t+1} + \beta \varphi_\mu \frac{\sigma_{\nu,t}^2}{\sigma_{\nu,t+1}^2 + \sigma_{\epsilon,t}^2} \mu_{\epsilon,t} \\ &+ \frac{\beta}{\alpha} \log \left(\mathbb{E}_t \exp \left[\alpha \left(\varphi_{M_1} - \varphi_{M_2} + \varphi_{\mu,t} \frac{\sigma_{\epsilon,t}^2}{\sigma_{\nu,t+1}^2 + \sigma_{\epsilon,t}^2} \right) (\epsilon_t + \nu_t) \right] \right). \end{aligned} \quad (43)$$

Matching the coefficients of the informational state $\mu_{\epsilon,t}$ to make the Bellman equation independent of its level delivers equation (17) for the shadow value φ_μ . Solving inductively the remaining state-independent terms in equation (43) for the welfare difference between the

²⁶In 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.

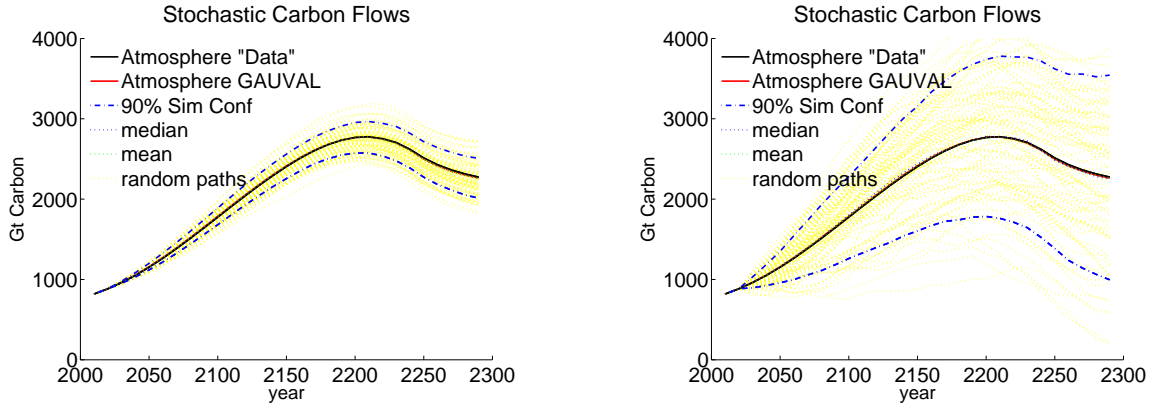


Figure 7: shows the evolution of atmospheric carbon under the low and the high specifications of the carbon cycle shock in equation (13), $\sigma_\chi = 10$ Gt on the left and $\sigma_\chi = 50$ Gt on the right. The shock's persistence of $\gamma_M = 0.997$ is calibrated to Joos et al.'s (2013) model 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.

uncertain and the deterministic scenario as in the proof of Proposition 4 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 (17). Canceling terms in the numerator of the expression in squared brackets delivers equation (18) in the main text.

E.3 Quantitative Analysis of Carbon Cycle Uncertainty

The quantification of carbon cycle uncertainty in section 4.4 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 4.2, 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 4.3, 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.

F Temperature Uncertainty: Details

F.1 Proof or Proposition 6

In the combined model of persistent epistemological and VAR uncertainty over the temperature increase in section 5.2 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 $\kappa_i, i \in \mathbb{N}$, eliminates the squared bracket in front of the cumulants and delivers the shadow values stated in equation (22). The difference between the uncertain and the deterministic value function's affine components derives analogously to the proof of Proposition 4 to

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

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 (21) in the main text.




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.

F.2 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 (19), 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 5.3, 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).²⁷

²⁷The 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.