

NBER WORKING PAPER SERIES

OPTIMAL CARBON ABATEMENT IN A STOCHASTIC EQUILIBRIUM MODEL
WITH CLIMATE CHANGE

Christoph Hambel
Holger Kraft
Eduardo Schwartz

Working Paper 21044
<http://www.nber.org/papers/w21044>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2015, Revised July 2017

We thank Christian Gollier, Frederick van der Ploeg, Armon Rezai, and Richard Tol for helpful comments and suggestions. We also thank participants of the 23rd Annual Conference of the European Association of Environmental and Resource Economists (EAERE 2017), the FinanceUC Conference, Santiago, Chile, the Santiago Finance Workshop, Chile, the Frankfurt-Mannheim Macro Workshop, and the joint seminar of the Humboldt University Berlin and ESMT for their comments and suggestions. All remaining errors are of course our own. Holger Kraft and Christoph Hambel gratefully acknowledge financial support by Deutsche Forschungsgemeinschaft (DFG), the Center of Excellence SAFE, funded by the State of Hessen initiative for research LOEWE. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Christoph Hambel, Holger Kraft, and Eduardo Schwartz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Carbon Abatement in a Stochastic Equilibrium Model with Climate Change
Christoph Hambel, Holger Kraft, and Eduardo Schwartz
NBER Working Paper No. 21044
March 2015, Revised July 2017
JEL No. D81,Q54

ABSTRACT

This paper studies a dynamic stochastic general equilibrium model involving climate change. Our framework allows for feedback effects on the temperature dynamics. We are able to match estimates of future temperature distributions provided in the fifth assessment report of the IPCC (2014). We compare two approaches to capture damaging effects of temperature on output (level vs. growth rate impact) and combine them with two degrees of severity of this damage (Nordhaus vs. Weitzman calibration). It turns out that the choice of the damage function is crucial to answer the question of how much the distinction between level and growth rate impact matters. The social cost of carbon is similar for frameworks with level or growth rate impact if the potential damages of global warming are moderate. On the other hand, they are more than twice as large for a growth rate impact if damages are presumably severe. We also study the effect of varying risk aversion and elasticity of intertemporal substitution on our results. If damages are moderate for high temperatures, risk aversion only matters when climate change has a level impact on output, but the effects are relatively small. By contrast, the elasticity of intertemporal substitution has a significant effect for both level and growth rate impact. If damages are potentially severe for high temperatures, then the results also become sensitive to risk aversion for both damage specifications.

Christoph Hambel
Goethe University
Department of Finance
Frankfurt
Germany
christoph.hambel@hof.uni-frankfurt.de

Eduardo Schwartz
Anderson Graduate School of Management
UCLA
110 Westwood Plaza
Los Angeles, CA 90095
and NBER
eduardo.schwartz@anderson.ucla.edu

Holger Kraft
Goethe University
Grueneburgplatz 1
65323 Frankfurt am Main
Germany
holgerkraft@finance.uni-frankfurt.de

1 Introduction

Our paper proposes a stochastic optimization-based general equilibrium model for the optimal abatement policy and optimal consumption. In contrast to most of the literature we allow for random evolutions of the key variables such as CO₂ concentration, global temperature and world GDP. We determine the optimal abatement policy and study this policy across different future scenarios for several model specifications. We provide detailed calibrations where we simultaneously match two decisive climate-sensitivity measures (TCR, ECS) which play an important role in the report of the IPCC (2014). In particular, we analyze the implications of different assumptions about the impact of climate change on output (level or growth) if damages of global warming are more or less pronounced (Nordhaus and Sztorc (2013) or Weitzman (2012)). We identify situations where the difference between growth rate and level impact matters and also study the effect of risk aversion and elasticity of intertemporal substitution (EIS).

If climate damages are moderate as in Nordhaus and Sztorc (2013), median results such as size of the social cost of carbon (SCC) over the next 100 years are similar for level and growth rate impact. However, a level impact induces a high variation in optimal emissions, whereas a growth rate impact leads to a less state-dependent abatement policy. The risk aversion has a smaller effect than the elasticity of intertemporal substitution (EIS) and is negligible for the growth rate impact. If climate damages are severe for high temperatures as in Weitzman (2012), there are pronounced differences between level and growth rate impact. Now, a growth rate impact leads to significantly higher SCC and induces a more state-dependent abatement policy. Furthermore, the results are sensitive to both the choice of risk aversion and EIS. In any case, abatement activities are more drastic, and optimal emissions are more volatile.

Notice that the social cost of carbon – an indicator of damage done by emitting carbon dioxide – is tightly varying with the optimal abatement policy. It is thus a stochastic process that increases in temperature and that depends on risk aversion and EIS in a similar way as the optimal abatement policy. Depending on the specification the optimal policy may or may not vary significantly with temperature. For instance, for a growth rate effect the policies are less state-dependent for moderate damages, but are crucially state-dependent for severe damages. Therefore, predications of deterministic models (e.g. DICE model) are also more or less accurate depending on the true nature of the effect. This is however an empirical question that is beyond the scope of our study.

Our paper is related to several studies using integrated assessment models (IAM):¹ Firstly, the DICE model (Dynamic Integrated Model of Climate and the Economy) is the most common framework to study optimal carbon abatement. It combines a Ramsey-type model for capital

¹IAMs can broadly be divided into two classes: *welfare optimization models* which choose an optimal abatement policy and *simulation models* that renounce an optimization routine and rather evaluate specific policy scenarios. Such a framework combines knowledge from different areas of science to an unified model that describes interactions between greenhouse gas emissions, the climate system and the economy.

allocation with deterministic dynamics of emissions, carbon dioxide and global temperature. Notice that it is formulated in a deterministic setting, see for example Nordhaus (1992, 2008), Nordhaus and Sztorc (2013).² Kelly and Kolstad (1999) and Kelly and Tan (2015) extend this model and allow the decision maker to learn about the unknown relation between greenhouse gas emissions and temperature. In frameworks with recursive utility, Crost and Traeger (2014), Jensen and Traeger (2014), and Ackerman et al. (2013) analyze versions where one component is assumed to be stochastic. In contrast to our paper, Crost and Traeger (2014) and Jensen and Traeger (2014) do not allow for stochastic temperature dynamics, but consider uncertainties in economic growth and the damage function. Ackerman et al. (2013) introduce transitory uncertainty of the climate sensitivity parameter into the DICE model. These studies indicate that risk aversion has a smaller effect on the social cost of carbon than the elasticity of intertemporal substitution. We confirm the earlier results for moderate climate damages and show that risk aversion can have a crucial effect if damages are severe for high temperatures. Cai et al. (2015) study a stochastic generalization referred to as DSICE model. Their approach is computationally involved since it is based on high-dimensional Markov chains. However, both carbon and temperature dynamics are deterministic and the model only involves a level impact on economic growth induced by climate change (as all variants of the DICE model).

By contrast, we suggest a flexible IAM that does not fall into the class of DICE models.³ All components are genuinely stochastic. In particular, it involves parsimonious stochastic dynamics for the world temperature. We can thus offer a way to simultaneously calibrate two decisive climate-sensitivity measures (TCR, ECS). Transient climate response (TCR) measures the total increase in average global temperature at the date of carbon dioxide doubling. Equilibrium climate sensitivity (ECS) refers to the change in global temperature that would result from a sustained doubling of the atmospheric carbon dioxide concentration after the climate system will have found its new equilibrium. Notice that the fifth assessment report of the IPCC (2014) provides estimates for the distribution of these measures in the future. In particular, it gives empirical evidence about non-negligible risks for extreme climate changes. In other words, future temperatures are supposed to be fat-tailed and right-skewed. Our approach allows us to match moments beyond the first and second moment. This gives us the opportunity to study the effects of fat-tailed and skewed temperature distributions and to capture some of the inherent uncertainty of the problem.⁴

Other related literature includes Bansal and Ochoa (2011), Dell et al. (2009, 2012) and Burke et al. (2015) who provide empirical evidence that temperature negatively affects the growth rate of output rather than its level (as in the DICE approach). Pindyck (2011, 2014) studies the effect

²When we refer to DICE in this article, we mean the DICE-2013R-version that is presented in Nordhaus and Sztorc (2013).

³To simplify comparisons with the existing literature, we calibrate the growth rate in the benchmark case such that our economic model closely matches the business-as-usual (BAU) evolution of GDP in DICE.

⁴See, e.g., the remarks of Nordhaus (2008) on the uncertainty of the problem.

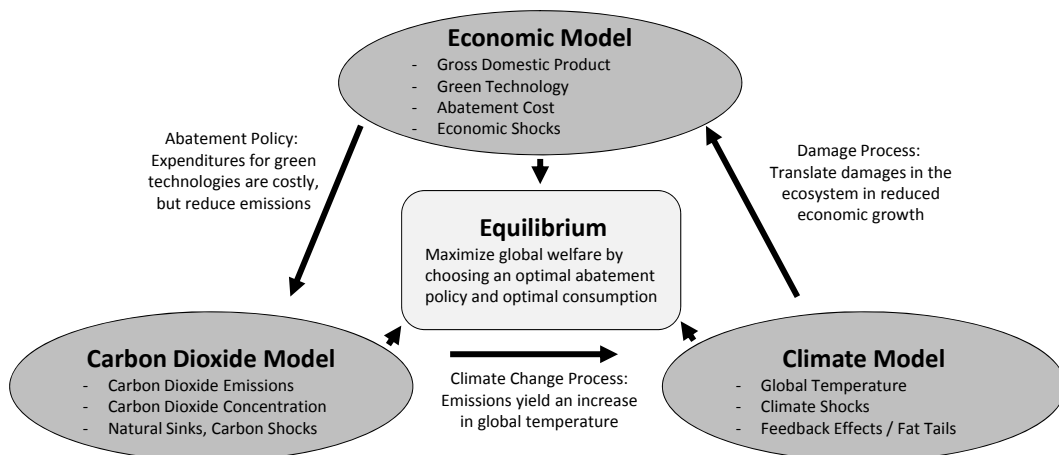


Figure 1: Building Blocks of the Model.

of a growth rate impact in an endowment economy. He solves a static instead of a dynamic optimization problem and calculates the so-called willingness to pay. This is the fraction of consumption that is necessary to keep global warming below some target temperature, e.g. 3°C . However, he abstracts from carbon dioxide emissions and abatement costs. Dietz and Stern (2015) study a stochastic version of DICE that is plagued by persistent impacts on economic growth and involves a fat-tailed ECS. Using a Monte-Carlo approach, they provide a solution where decisions are formed before the first period and are not revised. Moore and Diaz (2015) study the effect of growth rates impacts in a deterministic two-region version of DICE and find that a growth rate impact warrants stringent mitigation policy. Finally, similar as in our paper, Pindyck (2012) studies the difference between level and growth rate effect, but in the stylized setting of Pindyck (2011, 2014).

As in most of the above-mentioned papers, the starting point for our economic analysis of climate change is an integrated assessment model. Consequently, our model consists of three components: carbon dioxide model, climate model, and economic model. Section 2 describes these components and characterizes the equilibrium of the economy. Section 3 calibrates all model components. Section 4 presents our benchmark results. Additional robustness checks can be found in Section 5. Section 6 concludes. An Appendix contains additional material.

2 Model Setup

This section presents the model setup and describes its equilibrium. Figure 1 depicts the three building blocks of our framework (carbon dioxide model, climate model, and economic model).

The *carbon dioxide model* keeps track of the carbon dioxide concentration in the atmosphere. This concentration increases by anthropological and also non-man made carbon dioxide shocks

and it decreases since natural sinks such as oceans absorb carbon dioxide. Society can control anthropological carbon dioxide emissions by choosing an abatement strategy which reduces the current (business-as-usual) emissions. These efforts are costly.

The *climate model* measures the average world temperature and its departure from the pre-industrial level. Empirically, there is a (noisy) positive relation between carbon dioxide concentration and world temperature. Our temperature process captures this relation and allows for possible feedback effects.

The *economic model* describes the dynamics of global GDP (syn. output) in a stylized production economy. In our benchmark setting, global warming has a negative influence on economic growth, i.e. on the drift of global GDP. Alternatively, we also study a framework with a level impact as in DICE. Society can only indirectly mitigate this damaging effect by choosing the above mentioned abatement strategy. This is the link of the economic model to the emission model, which completes the circle.

Society (syn. mankind or decision maker) chooses optimal consumption and an optimal abatement strategy whose costs contemporaneously reduce economic growth. The remaining part of output must be invested so that an *equilibrium* materializes.

2.1 Carbon Dioxide Model

The average pre-industrial concentration of carbon dioxide in the atmosphere is denoted by Y^{PI} . The total current concentration of carbon dioxide in the atmosphere is given by

$$Y_t^\Sigma = Y^{\text{PI}} + Y_t, \quad (1)$$

where Y_t denotes the amount of atmospheric carbon dioxide that is caused by human activities, i.e. the part of atmospheric carbon dioxide that exceeds the pre-industrial concentration. Its dynamics are assumed to be

$$dY_t = Y_t [(\mu_y(t) - \alpha_t)dt + \sigma_y dW_t^y]. \quad (2)$$

We refer to (2) as carbon dioxide dynamics or process. Here $W^y = (W_t^y)_{t \geq 0}$ is a standard Brownian motion that models *unexpected* shocks on the carbon dioxide concentration. These could be the result of environmental shocks such as volcano eruptions or earthquakes, but they can also be man-made. The volatility of these shocks σ_y is assumed to be constant. Atmospheric carbon dioxide increases with an *expected* growth rate of μ_y that models the current growth path of the carbon dioxide concentration. In other words, μ_y is the growth rate if society does not take additional actions to reduce carbon dioxide emissions. We thus refer to μ_y as the *business-as-usual drift* of the carbon dioxide process. Notice that it also involves all *past* policies which have been implemented to reduce carbon dioxide emissions.

Society can however pursue new policies to reduce emissions. We refer to such an additional effort as an *abatement strategy* $\alpha = (\alpha_t)_{t \geq 0}$. In other words, the abatement policy α models how *additional* actions change the expected growth of the carbon dioxide concentration, i.e. these are abatement policies beyond business-as-usual (BAU). By definition, this differential abatement policy has been zero in the past ($\alpha_t = 0$ for all $t < 0$). If no abatement policy is chosen and society sticks to BAU, we also use the notation Y^{BAU} instead of Y .

Our dynamics of the carbon concentration Y are formulated in terms of the abatement policy α . However, we are also interested in the resulting CO₂ emissions. To back out the implied CO₂ emissions that are consistent with (2), we now consider alternative dynamics of Y where – up to environmental shocks – the change in Y is expressed as the difference between CO₂ emissions and the amount of carbon absorbed by natural sinks. Formally, if ϵ_t denotes the time- t anthropological carbon dioxide emissions, then we obtain

$$dY_t = \zeta_\epsilon \epsilon_t^\alpha dt - \delta_y(X_t) Y_t dt + Y_t \sigma_y dW_t^y, \quad (3)$$

$$dX_t = \delta_y(X_t) Y_t dt, \quad (4)$$

where ζ_ϵ is a factor converting emissions into concentrations.⁵ The variable X_t measures the total quantity of atmospheric carbon dioxide that has already been absorbed by natural sinks. The function δ_y models the decay rate of carbon dioxide, i.e. the speed at which carbon dioxide is absorbed from the atmosphere. We assume δ_y to decrease in X , i.e. the capacity of natural sinks declines with the quantity of carbon that has already been absorbed. Equation (3) can be considered as an *ecological budget constraint*: The total change in carbon dioxide is (up to environmental shocks) the difference between anthropological emissions and natural carbon sequestration.

The dynamics (2) and (3) can be interpreted as a system of two equations with the two unknowns dY and ϵ^α . By equating (2) and (3), we can solve for the anthropological emissions of carbon dioxide (short: emissions) that are consistent with both dynamics:

$$\epsilon_t^\alpha = \frac{Y_t}{\zeta_\epsilon} [\mu_y(t) + \delta_y(X_t) - \alpha_t]. \quad (5)$$

Equation (5) provides the relation between the abatement strategy and the anthropological emissions under that strategy. We use the notation ϵ_t^{BAU} for business-as-usual emissions ($\alpha = 0$).

Finally, we define the so-called *emission control rate* as

$$\varepsilon^\alpha = (\epsilon^{\text{BAU}} - \epsilon^\alpha) / \epsilon^{\text{BAU}} = 1 - \frac{Y_t}{\zeta_\epsilon \epsilon_t^{\text{BAU}}} (\mu_y(t) + \delta_y(X_t) - \alpha_t). \quad (6)$$

⁵Carbon dioxide emissions are measured in gigatons (GtCO₂), whereas concentrations are measured in parts per million (ppm).

This quantity denotes the fraction of abated carbon dioxide emissions compared to BAU. Equivalently, it is the percentage of carbon dioxide emissions that is prevented from entering the atmosphere if the abatement policy α is implemented. As in the DICE model, we assume that the emission control rate ε^α is between 0 and 1. The assumption $\varepsilon^\alpha \geq 0$ excludes strategies that lead to emissions beyond BAU. On the other hand, the assumption $\varepsilon^\alpha \leq 1$ implies that emissions cannot be negative, which might only be possible if there are major technological breakthroughs (e.g. direct carbon removal (DCR)).

Notice that the restriction $\varepsilon^\alpha \leq 1$ yields to the following upper bound on the abatement policy

$$\alpha_t \leq \mu_y(t) + \delta_y(X_t) \quad (7)$$

i.e. technological restrictions prevent society from implementing very high abatement policies. This constraint makes it harder to make up for opportunities that have been missed in the past.⁶

2.2 Climate Model

We assume that the average global increase in temperature from its pre-industrial level is given by the dynamics⁷

$$dT_t = \frac{Y_t \eta_\tau}{Y_t^\Sigma} (\mu_y(t) - \alpha_t) dt + \frac{Y_t \sigma_\tau}{Y_t^\Sigma} \left(\rho_{y\tau} dW_t^y + \sqrt{1 - \rho_{y\tau}^2} dW_t^\tau \right) + \theta_\tau(T_{t-}) dN_t^\tau. \quad (8)$$

We refer to (8) as global warming process. The parameter η_τ is a constant relating the change in global temperature to changes in carbon dioxide concentration. The Brownian motions W^τ and W^y are independent. The diffusion parameter σ_τ is assumed to be constant. Furthermore, $N^\tau = (N_t^\tau)_{t \geq 0}$ is a self-exciting process whose jump intensity $\pi_\tau(T_t)$ and jump size $\theta_\tau(T_t)$ can depend on T_t itself. There is empirical evidence that the distribution of future temperature changes is right-skewed (see IPCC (2014)). One reason for this is that there might be delayed climate feedback loops triggered by increases in global temperature. This line of argument suggests that the temperature dynamics involve a self-exciting jump process whose jump intensity and jump size depend on the temperature itself. Intuitively, this means that an increase in temperature makes future increases both more likely and potentially more severe. Therefore, a self-exciting process captures the idea of feedback loops and at the same time allows for calibrating the skewness of the distribution of future temperature changes.

⁶If it were really possible to actively remove carbon dioxide from the atmosphere (direct carbon removal), then negative CO₂ emissions would be feasible. As in the DICE model, we do not allow for negative emissions in our benchmark calibration. However, our results are robust to this assumption. In robustness checks not reported here, we have assumed that the emission control rate is restricted to $\varepsilon_t^\alpha \in [0, 1.2]$, where $\varepsilon_t^\alpha > 1$ involves negative emissions. On a time scale of 100 years, our median main results however hardly change. Only on extreme paths, society implements more stringent abatement policies leading to negative emissions.

⁷Appendix A provides a motivation for these climate dynamics and describes the technical details.

2.3 Economic Model

2.3.1 GDP Growth

Following Pindyck and Wang (2013) we assume that real GDP (syn. output) is generated by a production technology that is linear in capital (AK -technology) and involves capital adjustment costs.⁸ In our model, global warming has a negative impact on GDP. We assume that real GDP is

$$C_t = \widehat{C}_t \mathcal{K}_t^\alpha D_t \quad (9)$$

where \mathcal{K}_t^α is a factor that captures the costs resulting from implementing the abatement policy α . Section 2.3.2 below describes how \mathcal{K}_t^α is related to the abatement policy. D_t is a damage factor that scales down output in response to climate change. Section 2.3.3 discusses the form of the damage factor D_t in greater detail. In particular, we study a damage factor that influences the growth rate of GDP. Notice that \widehat{C}_t denotes real GDP in the *absence* of global warming (*shadow GDP*). Put differently, \widehat{C}_t is what GDP would be if there were no climate change. Its dynamics are given by

$$d\widehat{C}_t = \widehat{C}_t [g(t, \chi_t)dt + \sigma_c(\rho_{cy}dW_t^y + \widehat{\rho}_{c\tau}dW_t^\tau + \widehat{\rho}_c dW_t^c)], \quad (10)$$

where $W^c = (W_t^c)_{t \geq 0}$ is a third Brownian motion that is independent of W^y , W^τ and N^τ . The volatility σ_c of the economic shocks is assumed to be constant. Output is correlated with carbon concentration and temperature via ρ_{cy} and $\rho_{c\tau}$. Standard arguments then lead to the following specifications:⁹

$$\widehat{\rho}_{c\tau} = \frac{\rho_{c\tau} - \rho_{cy}\rho_{y\tau}}{\sqrt{1 - \rho_{cy}^2}}, \quad \widehat{\rho}_c = \sqrt{1 - \rho_{cy}^2 - \widehat{\rho}_{c\tau}^2}.$$

Shadow GDP evolves with an expected growth rate $g(t, \chi)$. Here, χ denotes the *consumption rate*, i.e. the fraction of GDP that is consumed. In our framework, the expected growth rate of shadow GDP has the following form (see Appendix B for more details)

$$g(t, \chi) = A(1 - \chi) - \delta_k(t) - \frac{1}{2}\vartheta(1 - \chi)^2, \quad (11)$$

where the capital-output ratio A is constant (AK -technology). The deterministic function $\delta_k(t)$ describes the trend of the capital stock dynamics. It captures the effects of depreciation, but also the growth of human capital. The parameter ϑ measures the costs of installing new capital. This specification ensures that capital adjustment costs are homogenous in capital and convex

⁸Formally, output is $C_t = AK_t$, where $A > 0$ is a constant modeling productivity and K is the only factor of production. K is the total stock of capital, i.e. it includes physical capital, but also human capital and firm-based intangible capital. Appendix B describes how the GDP dynamics can be derived from such a production technology.

⁹Formally, this is a Cholesky decomposition.

in investment and abatement. Such an assumption is widely used in the literature (see, e.g., Hayashi (1982) and Jermann (1998), Pindyck and Wang (2013)).

2.3.2 Abatement Costs

Output C can be used for consumption \mathcal{C} , investments \mathcal{I} or abatement expenditures \mathcal{A}^α ,

$$C_t = \mathcal{C}_t + \mathcal{I}_t + \mathcal{A}_t^\alpha. \quad (12)$$

Implementing an abatement policy α can be considered as spending money on more carbon efficient technologies that reduce carbon dioxide emissions (*green technologies*). The expenditures necessary to implement the policy α are the so-called abatement expenditures denoted by \mathcal{A}_t^α . To keep the model tractable, we measure and calibrate these expenditures relative to output.¹⁰ As can be seen from (9), we model the negative effects of an abatement policy α on output via a *cost process* \mathcal{K}_t^α that scales down output as a response to abatement expenditures. We assume that the cost process \mathcal{K}_t^α is of finite variation and can be represented by a sufficiently smooth cost function κ that depends on time and the emission control rate:

$$\mathcal{K}_t^\alpha = \exp\left(-\int_0^t \kappa(s, \varepsilon_s^\alpha) ds\right). \quad (13)$$

Therefore, implementing α reduces economic growth over the small interval $[t, t+dt]$ by $\kappa(t, \varepsilon_t^\alpha)dt$. We refer to $\kappa(t, \varepsilon_t^\alpha)$ as the instantaneous growth effect of the abatement expenditures \mathcal{A}^α . Put differently, $\kappa(t, \varepsilon_t^\alpha)$ models the forgone economic growth resulting from spending money on abatement. Appendix B shows that $\kappa(t, \varepsilon_t^\alpha)$ can be calculated if the abatement expenditures \mathcal{A}^α , productivity A , and output C are known:

$$\kappa(t, \varepsilon_t^\alpha) = A \frac{\mathcal{A}_t^\alpha}{C_t}. \quad (14)$$

Using relation (6), we can rewrite κ in terms of time t , carbon concentration Y and abatement policy α . Therefore, we also use the notation $\kappa(t, Y_t, \alpha_t)$ instead of $\kappa(t, \varepsilon_t^\alpha)$.

2.3.3 Impact of Global Warming

This paper studies two approaches of how to model economic damages induced by climate change. First, we consider a widely used approach assuming that current temperatures directly affect the *level* of GDP (see, e.g., Nordhaus (2008), Tol (2002), Hope (2006)). Second, we analyze a framework that models damages as a negative effect on the *growth rate* of GDP, which is suggested by empirical evidence.

¹⁰This is standard in the literature (e.g. DICE). Relaxing this assumption would lead to an additional state variable.

Level Impact The standard approach assumes that global warming has a direct impact on the level of GDP which is captured by a sufficiently smooth damage function $D(T)$ with $D(0) = 1$, $\lim_{T \rightarrow \infty} D(T) = 0$. Here T denotes the temperature anomaly. Thus, real GDP at time t is

$$C_t = \widehat{C}_t \mathcal{K}_t^\alpha D(T_t). \quad (15)$$

Notice that the damage function only depends on the *current* temperature anomaly. Therefore, in this specification only the current level of GDP is affected, but not its growth rate. We thus refer to such a damage specification as level impact. As Pindyck (2012, 2013) points out, a level impact has the unrealistic drawback that damages from global warming are reversible. For instance, if global temperatures increase heavily over 100 years, but then come back to the present level, GDP would fully recover without any permanent losses. Therefore, the impact of global warming as modeled by (15) is transitory.

Growth Rate Impact There is a lot of empirical evidence that higher temperatures affect growth rates of GDP rather than just its level (see, e.g., Dell et al. (2009, 2012)). In line with these findings, Pindyck (2012, 2013) argues that many damages induced by climate change are likely to be permanent, rather than transitory (e.g. destruction of ecosystems, sea level rise, or heat-related fatalities). Therefore, it is relevant to study the effects of a growth rate impact as well. In this case, we assume that real GDP at time t is given by

$$C_t = \widehat{C}_t \mathcal{K}_t^\alpha \exp\left(-\zeta_d \int_0^t T_s^k ds\right) \quad (16)$$

where ζ_d and k are positive parameters that relate temperature increase to loss of economic growth. Apparently, in such a framework the growth rate depends on the whole path of the temperature anomaly rather than just on the current value as in (15). It thus captures long-lasting damages from climate change that are difficult or impossible to reverse (at least in human timescales). Combining (10) and (16) yields the following dynamics of real GDP

$$dC_t = C_t \left[(g(t, \chi) - \zeta_d T_t^k) dt + \sigma_c (\rho_{cy} dW_t^y + \widehat{\rho}_{c\tau} dW_t^\tau + \widehat{\rho}_c dW_t^c) - \kappa(t, \varepsilon_t^\alpha) dt \right], \quad (17)$$

i.e. the expected growth rate $g(t, \chi) - \zeta_d T_t^k$ is negatively affected by current temperatures. The dynamics (17) are the reason why we refer to this damage specification as growth rate impact.¹¹

2.4 Equilibrium

It is well-known that for a decision maker with CRRA utility changing the degree of relative risk aversion has at first sight a counterintuitive effect: The abatement policy is *less* stringent if risk

¹¹Notice that for both specifications the distribution of future climate damages is known. Therefore, the Dismal Theorem of Weitzman (2009) does not apply.

aversion increases.¹² In order to resolve this puzzle and to disentangle relative risk aversion from elasticity of intertemporal substitution, we follow Crost and Traeger (2014), Jensen and Traeger (2014) and Ackerman et al. (2013) and assume the decision maker's preferences to be of Epstein-Zin type. This allows us to analyze the effects of varying EIS and risk aversion separately. The society's time- t utility index $J_t^{\alpha,\chi}$ associated with a given abatement-consumption strategy (α, χ) over the infinite planning horizon $[0, \infty)$ is thus recursively defined by

$$J_t^{\alpha,\chi} = \mathbb{E}_t \left[\int_t^\infty f(C_s^{\alpha,\chi}, J_s^{\alpha,\chi}) ds \right], \quad (18)$$

where $C^{\alpha,\chi} = \chi C^{\alpha,\chi}$ denotes consumption. Following Duffie and Epstein (1992) the aggregator function f is given by the continuous-time Epstein-Zin aggregator

$$f(C, J) = \begin{cases} \delta \theta J \left[\left(\frac{C}{[(1-\gamma)J]^{\frac{1}{1-\gamma}}} \right)^{1-\frac{1}{\psi}} - 1 \right], & \psi \neq 1 \\ \delta(1-\gamma)J \log \left(\frac{C}{[(1-\gamma)J]^{\frac{1}{1-\gamma}}} \right), & \psi = 1 \end{cases} \quad (19)$$

with $\theta = \frac{1-\gamma}{1-1/\psi}$. The parameter $\gamma > 1$ measures the degree of relative risk aversion, $\psi > 0$ reflects the elasticity of intertemporal substitution (EIS), and $\delta > 0$ denotes the time-preference rate.¹³ For $\theta = 1$ (or equivalently $\psi = 1/\gamma$), the preferences simplify to standard time-additive CRRA utility with utility function $u(c) = \frac{1}{1-\gamma} c^{1-\gamma}$. For $\theta < 1$ (i.e. $\psi > 1/\gamma$) the agent prefers early resolution of uncertainty and is eager to learn outcomes of random events before they occur. On the other hand, if $\theta > 1$ (i.e. $\psi < 1/\gamma$) the agent prefers late resolution of uncertainty. Notice that although recursive utility allows to disentangle risk aversion from EIS, it does not allow to disentangle prudence from the other two parameters as well. Following Kimball and Weil (2009) prudence is given by $\gamma(1 + \psi)$. Therefore, risk aversion and EIS affect prudence in a linear way. We will discuss the impact of prudence in the robustness section where we vary risk aversion and EIS separately.

The decision maker chooses an admissible abatement-consumption strategy (α, χ) in order to maximize his utility index $J_t^{\alpha,\chi}$ at any point in time $t \in [0, \infty)$. An admissible strategy must ensure that output, consumption, investment and abatement expenditures remain positive, i.e. $C_t, \mathcal{C}_t, \mathcal{I}_t, \mathcal{A}_t^\alpha \geq 0$ for all $t \geq 0$. Furthermore, the abatement policy must satisfy (7) and lead to a positive emission control rate. The class of all admissible abatement-consumption strategies

¹²Pindyck (2013) explains this fact as follows: For a higher level of risk aversion, the marginal utility of consumption declines faster. However, consumption is expected to grow and consequently utility from future consumption decreases with risk aversion. For a higher level of risk aversion society thus implements a less stringent abatement policy leading to higher emissions and a higher global temperature.

¹³Although empirical evidence suggests that $\gamma > 1$ is the reasonable specification for the index of relative risk aversion, it is also possible to define aggregator functions for $\gamma \in [0, 1]$.

at time t is denoted by \mathfrak{A}_t . The indirect utility function is given by

$$J(t, c, x, y, \tau) = \sup_{(\alpha, \chi) \in \mathfrak{A}_t} \{J_t^{\alpha, \chi} \mid C_t = c, X_t = x, Y_t = y, T_t = \tau\} \quad (20)$$

We solve the utility maximization problem (20) by applying the dynamic programming principle. Details of the HJB equation and the solution method are presented in Appendix F.

2.5 Social Cost of Carbon

Our model can be used to calculate the social cost of carbon (SCC). Following Nordhaus and Sztorc (2013), Traeger (2014) and others, we define the social cost of carbon as the marginal rate of substitution between carbon dioxide emission and GDP. Formally, the social cost of carbon is given by

$$\text{SCC}_t = -\frac{\partial J_t}{\partial \mathbf{e}_t} / \frac{\partial J_t}{\partial C_t}. \quad (21)$$

Intuitively, the social cost of carbon measures the increase in current GDP that is required to compensate for economic damages caused by an marginal increase of time- t emissions. Therefore, SCC can be interpreted as an optimal *carbon tax*, i.e. the tax that compensates for the negative external effects from burning carbon. More details on how SCC is calculated can be found in Appendix F.

3 Calibration

This section provides a detailed calibration of all model components. Table 1 summarizes the calibration results and serves as our benchmark calibration. This calibration assumes a growth rate effect of climate change. We choose the year 2015 as starting point of our model ($t = 0$).¹⁴

3.1 Preferences

In order to disentangle risk aversion from elasticity of intertemporal substitution, we use recursive preferences. In the literature, there is no consensus on how to choose γ and ψ .¹⁵ Many studies that incorporate recursive utility in an IAM choose $\gamma = 10$ and ψ in the range between 0.5 and 1.5 (see, e.g., Ackerman et al. (2013), Crost and Traeger (2014), Jensen and Traeger (2014) and Cai et al. (2015)). We follow that literature and choose $\psi = 1$ as the benchmark value for EIS.

¹⁴Since DICE starts in 2010 and evolves in steps of 5 years, this assumption simplifies comparisons.

¹⁵Bansal and Yaron (2004) and Vissing-Joergensen and Attanasio (2003) combine equity and consumption data and estimate an EIS of 1.5 and a risk aversion in the range between 8 and 10. On the other hand, Hall (1988), Campbell (1999), Vissing-Joergensen (2002) estimate an EIS well below one.

Carbon Dioxide Model		
Y^{PI}	Pre-industrial carbon dioxide concentration	280
Y_0	Initial excess carbon dioxide concentration	121
ζ_c	Conversion factor	0.1278
σ_y	Carbon dioxide volatility	0.0078
Climate Model		
T_0	Current global warming	0.9
η_τ	Temperature scaling parameter	2.592
σ_τ	Temperature volatility	0.1
$\rho_{y\tau}$	CO ₂ /temperature correlation	0.04
Economic Model		
C_0	Initial GDP (trillion US-\$)	75.8
A	Productivity	0.113
ϑ	Adjustment cost parameter	0.372
σ_c	GDP volatility	0.0162
$\rho_{c\tau}$	GDP/temperature correlation	0
ρ_{cy}	GDP/CO ₂ correlation	0.29
ζ_d	Damage scaling parameter	0.00026
Preferences		
δ	Time-preference rate	0.015
γ	Relative risk aversion	10
ψ	Elasticity of Intertemporal Substitution	1

Table 1: Benchmark Calibration. This table summarizes the parameters of the benchmark calibration which is described in Section 3.

The time-preference rate is $\delta = 0.015$, which is a standard assumption in the IAM literature (see, e.g., the recent version of the DICE model by Nordhaus and Sztorc (2013)). In robustness checks, we vary these parameters and study their effects on our results.

3.2 Carbon Dioxide Model

The fifth assessment report of the IPCC (2014) provides four stylized climate scenarios depending on the future evolution of greenhouse gas emissions referred to as representative concentration pathways (RCPs). The RCP 8.5 scenario is characterized by high CO₂ emissions where the atmospheric concentration is supposed to stabilize at a high level in the second half of the 23th century.¹⁶ Consequently, the RCP 8.5 data is well-suited to serve as the average BAU scenario for CO₂ emissions and concentrations. Notice that all RCPs are deterministic, i.e. they can only be used to calibrate averages. Therefore, we use historical data to estimate the randomness of the carbon dioxide concentration.

Carbon Dioxide Dynamics To calibrate (1) and (2), we fix the pre-industrial carbon dioxide concentration at $Y^{\text{PI}} = 280$ ppm, which is a common assumption in the literature. Furthermore, in the year 2015 ($t = 0$) the carbon dioxide concentration was 401 ppm, which implies $Y_0 = 121$

¹⁶The data is available at <http://tntcat.iiasa.ac.at/RcpDb>

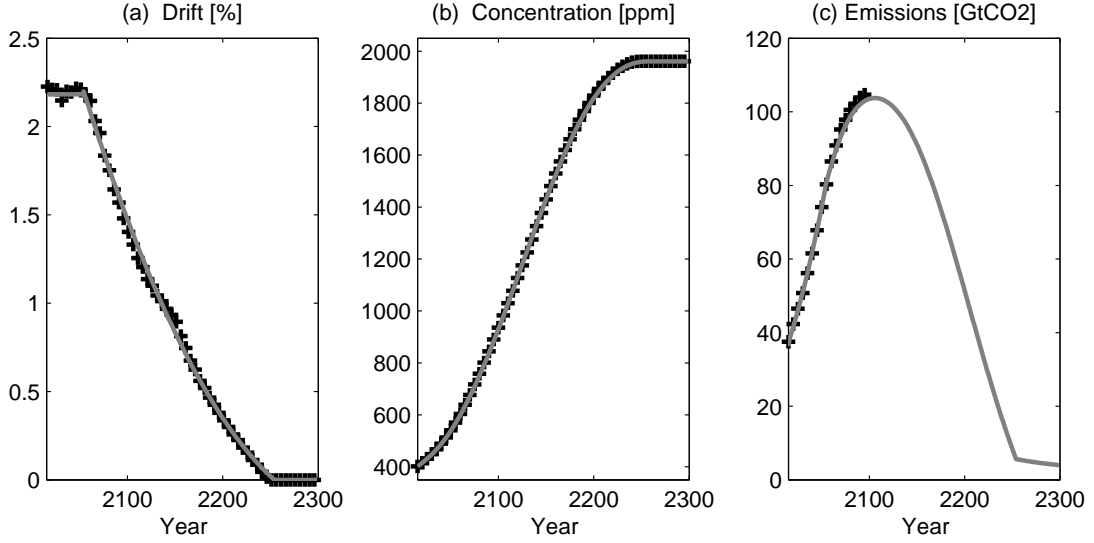


Figure 2: Calibration of the Carbon Dioxide Model. The crosses in Graph (a) depict the implied drift of the evolution of atmospheric carbon dioxide in the RCP 8.5 scenario. The solid line is our calibration of $\mu_y(t)$. The crosses in Graph (b) depict the evolution of atmospheric carbon dioxide in the RCP 8.5 scenario. The solid line shows our calibration to that data. The crosses in Graph (c) depict the emission prognosis in the RCP 8.5 scenario. The solid line shows our calibration to that data and an extension until 2300.

ppm as starting value for the carbon dioxide process (2). Then we calibrate the drift $\mu_y(t)$ such that the drift of the average BAU evolution (i.e. $\alpha = 0$ and $\sigma_y = 0$ in (2)) is close to the drift of the RCP 8.5 scenario that is marked by crosses in Graph (a) of Figure 2.¹⁷ Obviously, RCP 8.5 assumes three different regimes. For the first 40 years, the drift is virtually flat at a level close to the historical trend. Then the drift falls to zero over the next 200 years where it remains afterwards. This functional form of the drift rate can be captured in the following way:

$$\mu_y(t) = 0.022 \mathbf{1}_{\{t < 40\}} + (at^2 + bt + c) \mathbf{1}_{\{40 \leq t \leq 240\}} \quad (22)$$

where $a = 3.107 \cdot 10^{-7}$, $b = -1.963 \cdot 10^{-4}$, $c = 0.0292$. Graph (b) shows that, by applying (22), our median path simulated using the calibration of $\mu_y(t)$ (solid line) fits the the RCP 8.5 concentration data points (crosses) very well ($R^2 < 99\%$). To determine the volatility of carbon dioxide, we cannot apply the RCP 8.5 data which is deterministic. We thus use historical carbon dioxide records to estimate this parameter.¹⁸ Calculating the standard deviation of the log changes of Y yields a volatility of $\sigma_y = 0.0078$.

¹⁷We have calculated the drift of the RCP 8.5 scenario by computing the log-returns of the excess carbon dioxide concentration of two consecutive years.

¹⁸Source: Mauna Loa Observatory, Hawaii. Data available at <http://co2now.org/Current-CO2/CO2-Now/>.

Ecological Budget Constraint In a second step, we calibrate the decay rate of carbon dioxide $\delta_y(X_t)$ such that the model-implied carbon dioxide emissions (5) match the RCP 8.5 emissions (crosses in Graph (c) of Figure 2). The main issue here is that RCP 8.5 provides concentration data until 2300, but emission data only until 2100. We thus perform our calibration in two steps: First, we use both concentration and emission data until 2100 and determine the functional form of δ_y . Here we fix the conversion factor at $\zeta_c = 0.1278 \frac{\text{ppm}}{\text{GtCO}_2}$, which converts emissions into concentrations (see, e.g., IPCC (2014) and the references therein). Then we use this functional form and the concentration data to extrapolate the emissions until 2300.

As can be seen from Graph (b), the concentration of RCP 8.5 has an inflection point around 2100 and remains flat after the year 2240. Consequently, the emissions of RCP 8.5 must be hump-shaped. Since these emissions level off around 2100 in the data (crosses in Graph (c)), it is reasonable to expect a turning point around that year or shortly after, although - as noted above - RCP 8.5 is silent about emissions after the year 2100.¹⁹ This is exactly what our extrapolation yields.

The solid line in Graph (c) depicts the fit to that data and our BAU-emission forecast until 2300. It turns out that the following functional form of the decay rate of carbon dioxide matches the data well:

$$\delta_y(x) = a_\delta e^{-\left(\frac{x-b_\delta}{c_\delta}\right)^2}$$

where we estimate $a_\delta = 0.0176$, $b_\delta = -27.63$, $c_\delta = 314.8$ ($R^2 > 99\%$). Appendix C describes the technical details. Notice that the presumed evolution of BAU emissions beyond 2100 is similar to the baseline evolution in DICE. For instance, in the year 2200 DICE predicts 59GtCO₂, which is close to the estimate of 54GtCO₂ in our model.

3.3 Climate Model

The calibration of the global warming process (8) is divided into two steps: First, we calibrate the direct impact of the carbon dioxide concentration on global warming (captured by the continuous part of the model). The drift is calibrated using historical data, whereas the estimate of the volatility involves data on the transient climate response (TCR). In a second step, we calibrate the jump size and jump intensity such that the model can generate the above mentioned feedback effects. Here we use data on the equilibrium climate sensitivity (ECS).

Direct Impact: Drift and Volatility To estimate the parameter η_τ in the drift of the process, we use historical data on carbon dioxide concentration and global warming.²⁰ Notice

¹⁹Therefore, we can merely extrapolate the emissions from 2100 onwards. It is however obvious that concentration can only flatten out if emissions eventually decrease and reach a low level where natural sinks can absorb all emissions such that concentration does not increase any more.

²⁰Source: United Kingdom's national weather service. Global warming data available at <http://www.metoffice.gov.uk/>.

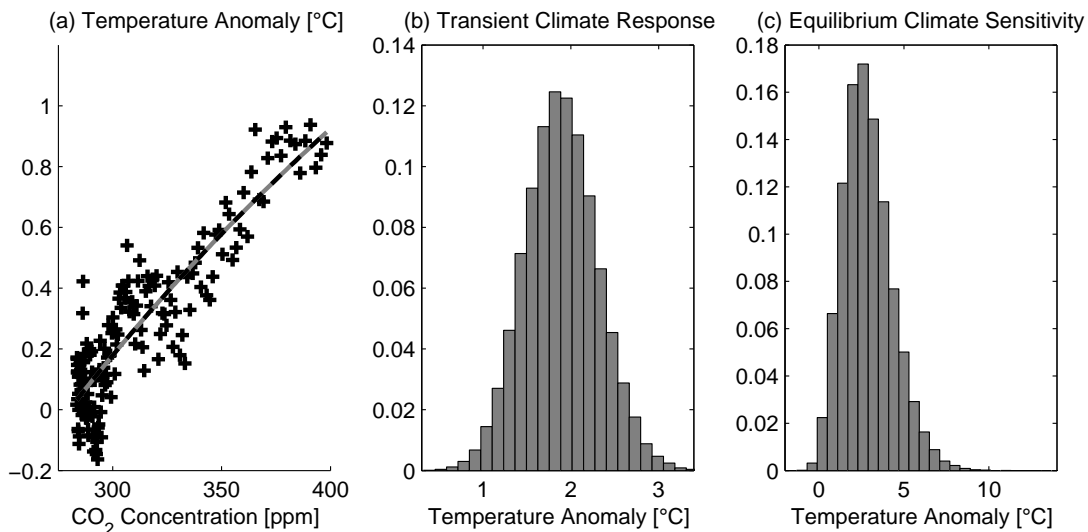


Figure 3: Calibration of the Climate Model. The crosses in Graph (a) depict pairs of empirical global warming and atmospheric carbon dioxide concentration. The solid line depicts the regression curve (23). The estimated parameters of the fitted curve is $\eta_\tau = 2.592$. Graph (b) shows a histogram of the simulated transient climate response. Graph (c) depicts a histogram of the equilibrium climate sensitivity. The histograms are based on a simulation of 1 million sample paths.

that the starting point for our model of the global warming dynamics is (24) in Appendix A. Therefore, we estimate η_τ by running a linear regression of global warming data on log-carbon dioxide data. Put differently, we calculate

$$\min_{\eta_\tau} \sum_{i=1}^N \left[T_i - \eta_\tau \log \left(\frac{Y_i^\Sigma}{Y^{\text{PI}}} \right) \right]^2. \quad (23)$$

Here T_i denotes the temperature above the pre-industrial level and Y_i^Σ denotes the carbon dioxide concentration at time t_i . Our estimation yields $\eta_\tau = 2.592$. The linear model performs well with $R^2 > 0.8$. Graph (a) of Figure 3 depicts the data and the estimate. We also use that data in order to estimate the correlation between carbon dioxide and global warming. We obtain a correlation parameter $\rho_{y\tau} = 0.04$.

To calibrate the diffusion coefficient σ_τ of (8), we use data on a measure called the *transient climate response* (TCR). TCR measures the total increase in average global temperature at the date of carbon dioxide doubling, $t_{2\times} = \inf_t \{t \geq 0 \mid Y_t = Y^{\text{PI}}\}$. The data comes from CMIP5.²¹ They simulate the future climate dynamics and obtain a multimodel mean (as well as median) of about $\mathbb{E}[\text{TCR}] = 1.8^\circ$ and a 90% confidence interval of $[1.2^\circ\text{C}, 2.4^\circ\text{C}]$. This points towards an approximately symmetric distribution of TCR, which is in line with our Brownian assumption. Further, notice that our above estimate of η_τ leads to a total temperature increase

²¹CMIP5 refers to Coupled Model Intercomparison Project Phase 5. See <http://cmip-pcmdi.llnl.gov/cmip5/> for further information.

of about $\eta_\tau \log(2) = 1.797$ at the relevant date $t_{2\times}$ for TCR. This is also in line with the CMIP5 estimate. Therefore, we are left with finding σ_τ , which we achieve by using the information about the confidence interval. The 95%-quantile is 1.65 standard deviations above the mean. This implies a standard deviation of $\sigma_{\text{TCR}} = 0.6^\circ\text{C}/1.65 = 0.364^\circ\text{C}$. We choose the volatility parameter σ_τ such that our model fits the distribution of TCR at the time when carbon dioxide is supposed to double. For this purpose, we estimate the doubling time $t_{2\times}$ via Monte Carlo simulation: We sample 1 million uncontrolled carbon dioxide paths to determine the time of carbon dioxide doubling. Then, we simulate 1 million global warming paths and choose the diffusion parameter such that the simulated distribution of TCR matches the above mentioned quantiles at the time of carbon dioxide doubling (see Graph (b) of Figure 3).²² On average, doubling occurs in 2055. As a result of the calibration, we estimate $\sigma_\tau = 0.1$ and a small correlation of about $\rho_{y\tau} = 0.04$.

Feedback Effects: Jumps In a second step, we calibrate the jump intensity and size using IPCC estimates for the *equilibrium climate sensitivity* (ECS). ECS refers to the change in global temperature that would result from a sustained doubling of the atmospheric carbon dioxide concentration after the climate system will have found its new equilibrium. This process is presumably affected by feedback effects kicking in *after* the temperature has increased significantly (e.g. the date related to TCR). Since the jump part in our model captures feedback effects, we use ECS data to estimate the corresponding parameters. Unfortunately, there is no consensus distribution for ECS because finding a new equilibrium might take hundreds of years. Summarizing more than 20 scientific studies, the IPCC (2014) however states that ECS is “likely” in the range of 1.5°C to 4.5°C with a most likely value of about 3°C .²³ Additionally, there is a probability of 5 to 10% that doubling the carbon dioxide concentration leads to an increase in global temperature of more than 6°C , while its extremely unlikely (i.e. less than 5%) that temperature increase is below 1°C . These numbers suggest that ECS has a right-skewed distribution which can be generated by jumps.

We assume that the climate system will find its new equilibrium 100 years after the carbon dioxide concentration will have doubled. We choose a functional form and an appropriate parametrization for the jump size and jump magnitude such that we can reproduce the above mentioned mean and quantiles of ECR by running Monte Carlo simulations. Furthermore, we perform the calibration in such a way that the constructed distribution for TCR is preserved. The latter is achieved by allowing for very small negative jumps when the temperature increase is still low. We thus choose the following parametrization of the climate shock intensity and

²²Here we set the jump part equal to zero such that the results are not driven by warming feedback effects. See also the definition of ECS in the next section.

²³In the language of IPCC, the word “likely” means with a probability higher than 67%.

Specification	Calibrated with respect to	Parametrization
Level Impact		
(L-N)	Nordhaus and Sztorc (2013)	$D^N(T) = \frac{1}{1+0.00266T^2}$
(L-W)	Weitzman (2012)	$D^W(T) = \frac{1}{1+(T/20.64)^2+(T/6.081)^{6.754}}$
Growth Impact		
(G-N)	Nordhaus and Sztorc (2013)	$\zeta_d^N = 0.00026, k = 1$
(G-W)	Weitzman (2012)	$\zeta_d^W = 0.000075, k = 3.25$

Table 2: Damage Specifications. The table summarizes the four different damage specifications that are studied in this paper.

magnitude:

$$\pi_\tau(\tau) = \left(\frac{0.95}{1 + 2.8e^{-0.3325\tau}} - 0.25 \right)^+,$$

$$\theta_\tau(\tau) = -0.0029\tau^2 + 0.0568\tau - 0.0577$$

Notice that we calibrate the jump intensity such that $\pi_\tau(\tau) = 0$ for all $\tau \leq 0$, i.e. there are no feedback effects if the global temperature is at or below its pre-industrial level. The simulated ECS distribution is depicted in Graph (c) of Figure 3.

3.4 Economic Model

We calibrate the growth rate (11) such that our economic model closely matches the evolution of GDP growth in the DICE model. Additionally, we chose the abatement cost function from DICE and derive the functional form of κ . The technical details can be found in Appendix D.

In order to analyze the impact of warming, we consider a set of possible specifications. The standard approach in the literature assumes that warming has a direct impact on the level of GDP via a sufficiently smooth damage function $D(T)$ with $D(0) = 1$. Thus, GDP at time t is $C_t = \widehat{C}_t D(T_t)$, where \widehat{C}_t denotes GDP in the absence of global warming (shadow GDP). There is however empirical evidence that rather the growth rate of GDP than the level is affected by global warming., e.g. Dell et al. (2009, 2012). To compare the effects of different damage types, we implement our model with four different specifications for the impact of warming. Table 2 summarizes these specifications.

Level Impact The standard damage function in DICE is inverse quadratic. Nordhaus and Sztorc (2013) use the parametrization

$$D^N(T) = \frac{1}{1 + 0.00266T^2},$$

which we refer to as (L-N) specification. They calibrate the damage function to temperature increases between 0°C to 3°C. They acknowledge that adjustments might be needed in case of higher warming. Weitzman (2012) proposes an alternative damage function that is based on an expert panel study involving 52 experts on climate economics. His damage function is designed to capture tipping point effects for very high temperature increases:

$$D^W(T) = \frac{1}{1 + (T/20.64)^2 + (T/6.081)^{6.754}},$$

which we refer to as (L-W) specification. The two damage functions are very close for temperatures in the range between 0°C and 3°C. From 3°C onwards, the losses start to deviate significantly. For instance, for a temperature increase of 6°C, Nordhaus' damage function D^N predicts a GDP loss of 9.2% percent, while Weitzman's specification D^W generates a loss of approximately 50% of GDP.

Growth Rate Impact To compare the effects of level and growth rate impacts, we first calibrate the growth rate impact such that the GDP dynamics are close to those resulting from a Nordhaus' level damage (L-N). In (17) we set $k = 1$. Furthermore, we choose the damage parameter $\zeta_d^N = 0.00026$ such that the average GDP losses in the year 2100 coincide for both specifications. Formally, using (15) and (16) the following equation implicitly determines ζ_d^N

$$\mathbb{E} \left[e^{-\zeta_d^N \int_0^t T_s ds + \sigma_c \widehat{W}_t^c} \right] = \mathbb{E} \left[e^{\sigma_c \widehat{W}_t^c} D^N(T_t) \right],$$

where $\widehat{W}_t^c = \rho_{cy} W_t^y + \widehat{\rho}_{c\tau} W_t^\tau + \widehat{\rho}_c W_t^c$ and t denotes the year 2100. We refer to the resulting specification as (G-N). Notice that this parameter is in line with the calibration of Pindyck (2014). Similarly, we calibrate the growth rate impact (G-W) such that the GDP dynamics are close to those resulting from a Weitzmans' level damage (L-W). This yields $k = 3.25$ and $\zeta_d^W = 0.000075$.

4 Main Results

This section presents our main results for the model introduced in Section 2. In particular, we determine the optimal abatement policy, its costs, the evolution of real GDP as well as the evolution of the carbon dioxide concentration and global average temperature changes over the next 100 years. Unless otherwise stated, we use our benchmark calibration from Section 3 that is summarized in Table 1.

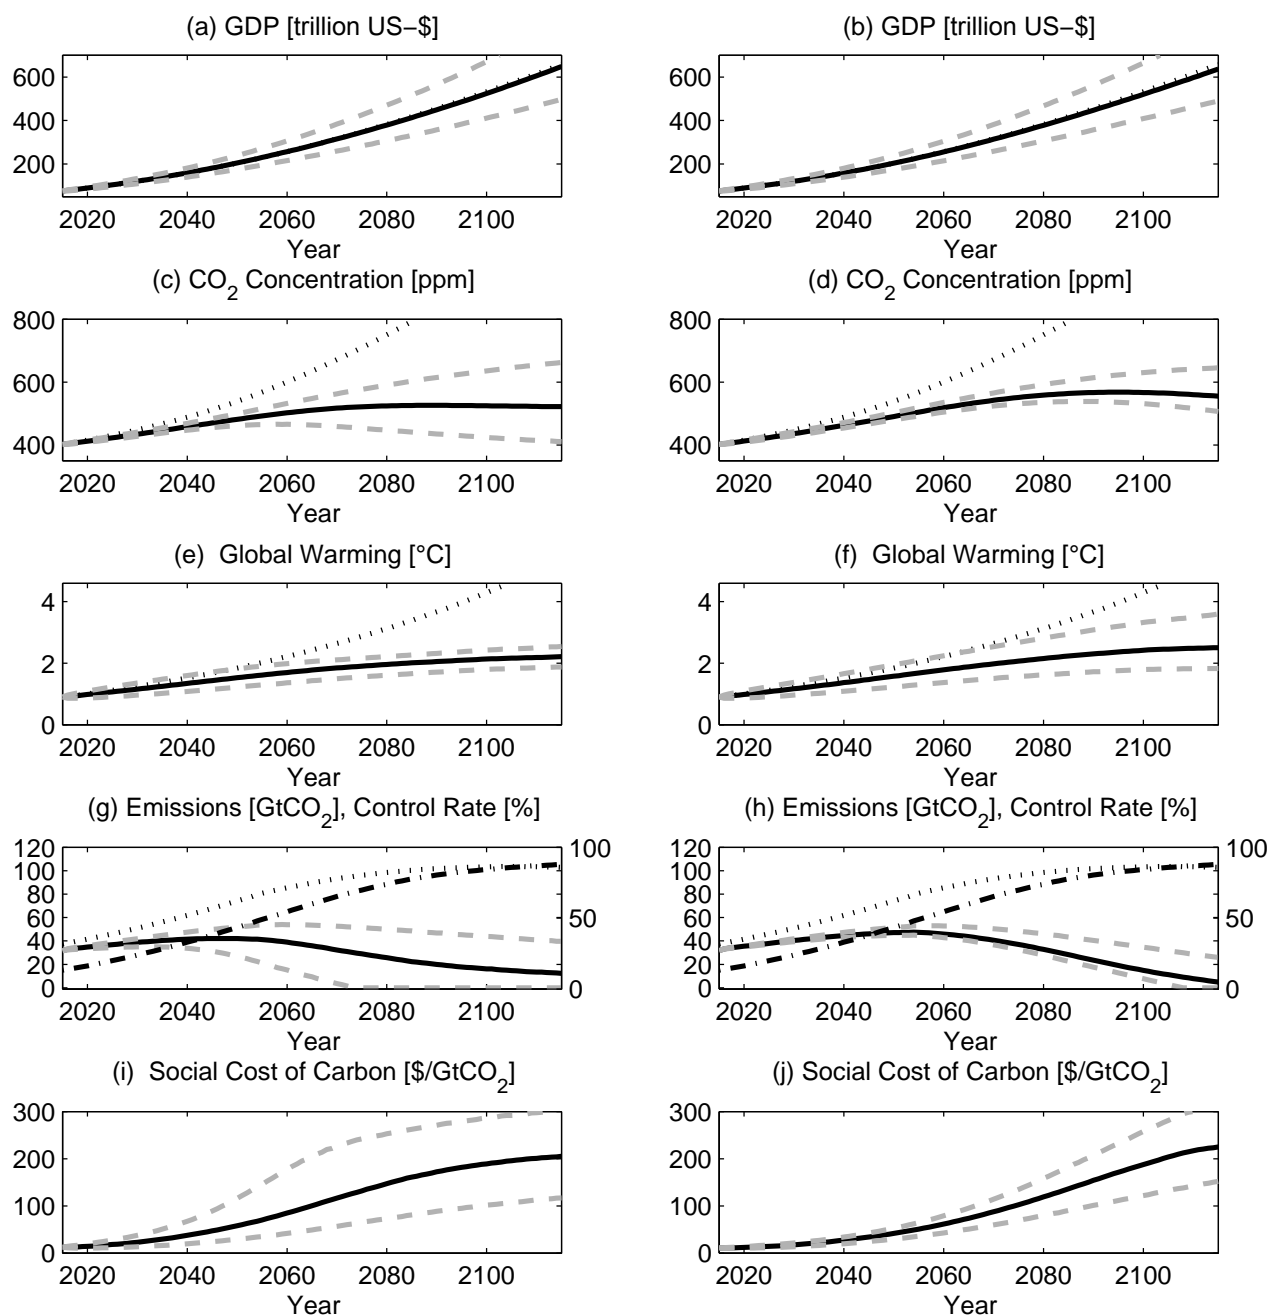


Figure 4: Results for the Nordhaus Calibration. Based on the calibration of Section 3, the graphs depict our results for the level impact (L-N) (left column) and the growth rate impact (G-N) (right column). Optimal paths are depicted by solid lines and BAU paths by dotted lines. Dashed lines show 5% and 95% quantiles of the optimal solution. Graphs (a) and (b) depict the evolution of world GDP, (c) and (d) the carbon dioxide concentration, (e) and (f) changes in global temperature, (g) and (h) carbon dioxide emissions and the median optimal emission control rate (dash-dotted line), (i) and (j) the social cost of carbon.

Model		2015	2035	2055	2075	2095	2115	2150	2200
(G-N)	GDP [trillion \$]	75.8	138.9	228.9	345.4	483.7	637.1	941.8	1435.6
	SCC [\$ /tCO ₂]	11.12	21.75	50.67	102.52	171.21	225.10	254.12	353.25
	Abatement Expenditures [trillion \$]	0.01	0.11	0.59	2.02	4.58	6.87	7.11	5.72
	Emission Control Rate	0.12	0.24	0.41	0.61	0.82	0.95	1	1
	Temperature rise [°C]	0.9	1.3	1.7	2.1	2.4	2.5	2.7	3.1
(L-N)	GDP [trillion \$]	75.8	139.3	230.4	348.7	490.2	652.7	988.0	1558.3
	SCC [\$ /tCO ₂]	10.63	24.23	58.37	116.84	183.03	221.77	254.70	376.68
	Abatement Expenditures [trillion \$]	0.01	0.12	0.73	2.41	4.98	6.69	6.99	6.00
	Emission Control Rate	0.12	0.25	0.44	0.65	0.83	0.93	1	1
	Temperature rise [°C]	0.9	1.3	1.7	2.0	2.3	2.4	2.5	2.9

Table 3: Median Results for the Nordhaus Calibration. The table reports the median evolution of selected variables for the growth rate (G-N) and level (L-N) impact.

4.1 Level vs. Growth Rate Impact for the Nordhaus Calibration

Table 3 and Figure 4 compare the evolutions of key state variables for the growth and level damage specifications (G-N) and (L-N). Both models behave similarly until the end of this century. This is not surprising as we calibrate the growth rate impact (G-N) such that the BAU evolution of world GDP until 2100 it is close to the one in (L-N). However, there are two main differences between these specifications. First, although the optimally controlled outputs in models (G-N) and (L-N) are similar until 2115, they diverge significantly in later years such as 2200 where the median output is almost 9% smaller in the model with growth impact. Second, the variation of global temperature in (G-N) is much higher than in (L-N), while the variability of emissions and concentrations is lower.

Notice that for a level impact damages are directly related to the current temperature, whereas for a growth rate impact damages depend on the whole temperature path so that the weight of the current temperature is much smaller and the damaging effects of high temperatures are delayed. Therefore, the abatement policy in (L-N) is slightly more stringent on average, but far more stringent for high temperatures. This implies a higher variability in carbon dioxide emissions and concentrations, but a lower variability in global temperature compared to (G-N). It also leads to a higher median output for (L-N), since more rigorous abatement policies tend to avoid economic damages more effectively.

Our analysis confirms and extends the results in Pindyck (2012). He shows in a static model that the *willingness to pay*²⁴ for keeping global warming below a certain threshold is higher for level damages than for growth damages, a finding that is in line with our results. However, Pindyck (2012) also states that there are no substantial differences between the two models. Our findings challenge this conclusion. First, output levels are significantly different in the year 2200, which is reported in Table 3. Second, the optimal emission path depends strongly on both

²⁴The willingness to pay is defined as the percentage of output that society is willing to sacrifice to keep the temperatures below a specified threshold.

Model		2015	2035	2055	2075	2095	2115	2150	2200
(G-W)	GDP [trillion \$]	75.8	138.1	223.4	330.3	459.4	610.4	921.1	1451.1
	SCC [\$ / tCO ₂]	42.86	92.55	145.93	172.82	188.20	198.44	219.77	333.00
	Abatement Expenditures [trillion \$]	0.14	0.99	3.06	4.59	5.37	5.63	5.58	4.89
	Emission Control Rate	0.30	0.53	0.75	0.83	0.87	0.89	0.92	0.96
	Temperature rise [°C]	0.9	1.1	1.2	1.3	1.4	1.4	1.4	1.4
(L-W)	GDP [trillion \$]	75.8	139.2	229.4	343.4	480.6	642.0	975.3	1551.0
	SCC [\$ / tCO ₂]	18.07	42.40	93.19	152.92	189.00	207.09	222.51	346.33
	Abatement Expenditures [trillion \$]	0.03	0.29	1.50	3.69	5.28	5.93	5.46	4.94
	Emission Control Rate	0.16	0.35	0.57	0.76	0.85	0.89	0.89	0.93
	Temperature rise [°C]	0.9	1.2	1.5	1.7	1.8	1.9	2.0	2.4

Table 4: Median Results for the Weitzman Calibration. The table reports the median evolution of selected variables for the growth rate (G-W) and level (L-W) impact.

the current state of the climate system and the damage specification. For instance by 2095, the 95% quantile of temperature is 3.1 (2.6) °C in the model with growth (level) impact leading to optimal carbon dioxide emissions of 19 (0) GtCO₂. This implies that the choice of the damage specification (growth rate or level impact) can have a significant effect, in particular for extreme paths.

4.2 Level vs. Growth Rate Impact for the Weitzman Calibration

We now consider the specifications (L-W) and (G-W), which are described in Section 3.4. Table 4 and Figure 5 show our corresponding findings. The graphs of Figure 5 on the left-hand (right-hand) side depict the results for the level (growth rate) impact. To avoid the potentially severe consequences of global warming, society keeps temperature low and in a narrow confidence band, which can be seen in Graphs (e) and (f). In turn, this leads to abatement strategies that are more sensitive to changes in current temperature. Therefore, most of the damaging effects of climate change can potentially be avoided resulting in steady economic growth (see Graphs (a) and (b)). From the end of the century onwards, the BAU paths of GDP are significantly lower than the optimally controlled paths.

Although in both scenarios society acts more rigorously than in the previous case, there are quantitative differences between the level impact (L-W) and the growth rate impact (G-W) that are also qualitatively different from our previous results on the Nordhaus calibration. For the growth rate impact, SCC is initially 42.86 and thus more than twice as high as for the level impact where it is 18.07. This implies more rigorous abatement activities in (G-W) than in (L-W). Therefore, the temperature increase is significantly smaller. Surprisingly, now the growth rate impact involves a higher SCC. This can intuitively be explained by the attitude of an agent with recursive preferences towards changes in the drift of his endowment stream. The long-run risk literature (see, e.g., Bansal and Yaron (2004)) documents that this type of agents is very sensitive to persistent changes of the growth rate. Whereas in the Nordhaus calibration the

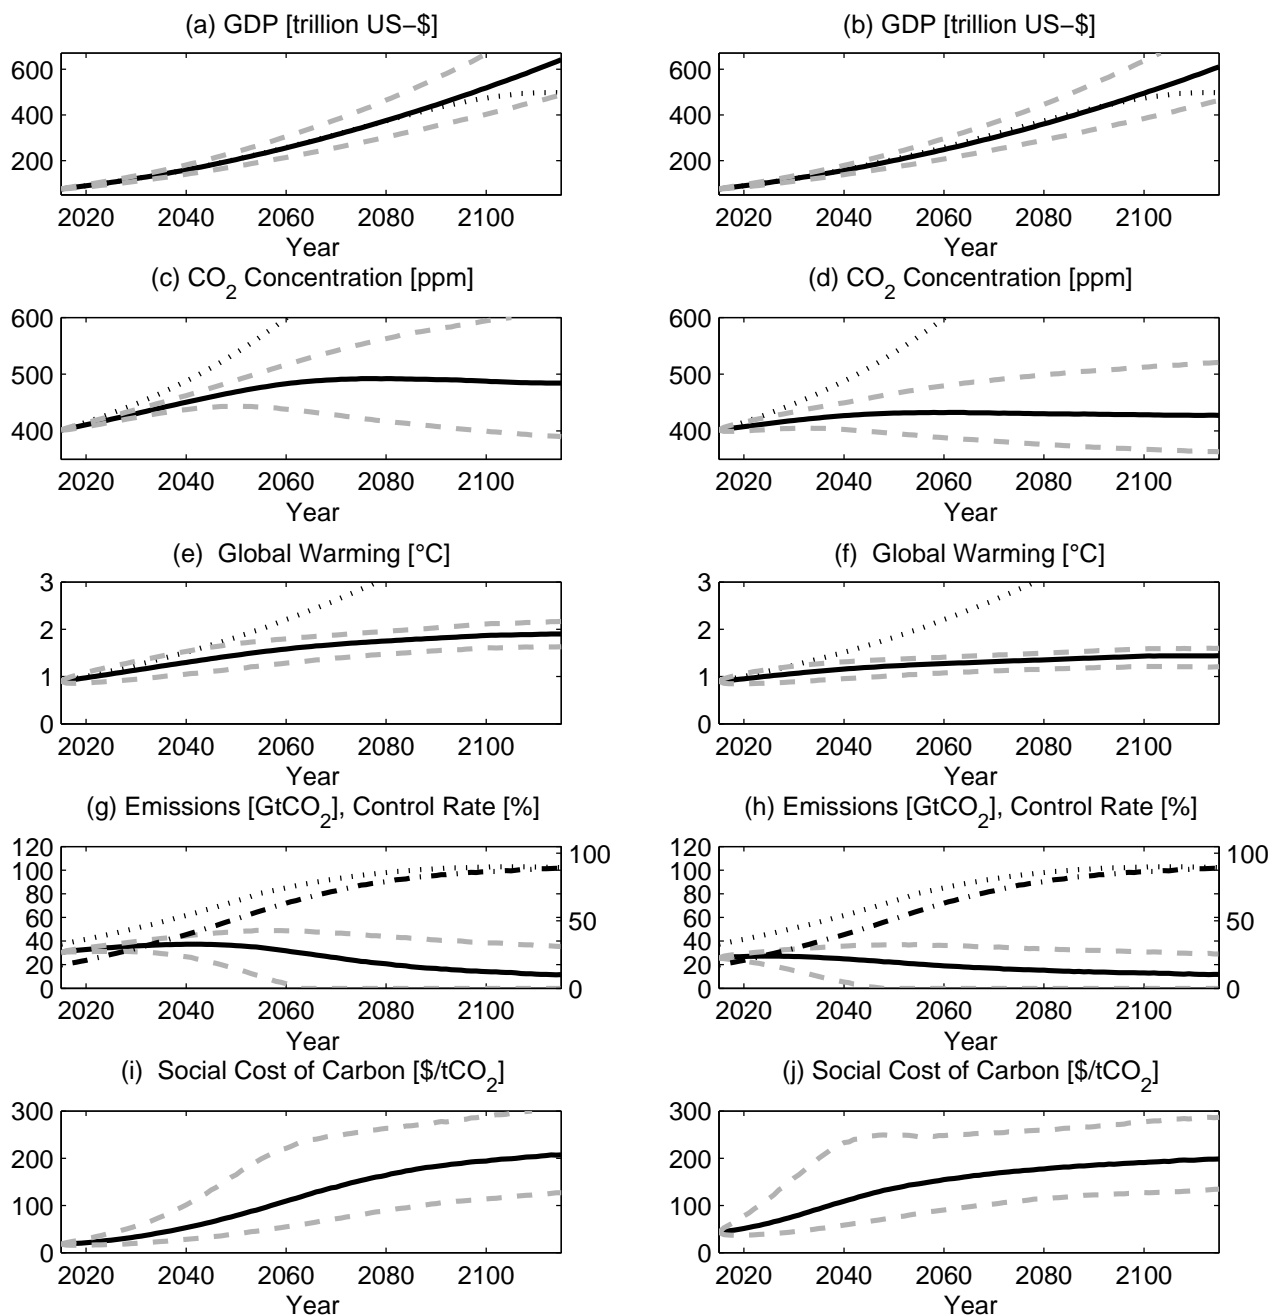


Figure 5: Weitzman Damage Specification. The graphs depict our results for the level impact (L-W) (left column) and the growth rate impact (G-W) (right column). Optimal paths are depicted by solid lines and BAU paths by dotted lines. Dashed lines show 5% and 95% quantiles of the optimal solution. Graphs (a) and (b) show the evolution of world GDP, (c) and (d) the carbon dioxide concentration in the atmosphere, (e) and (f) median changes in global temperature, (g) and (h) carbon dioxide emissions and the median optimal emission control rate (dash-dotted line), (i) and (j) the social cost of carbon.

ψ	γ				
	1	2	5	10	15
Nordhaus Calibration					
0.5	5.58 (5.75)	5.61 (5.80)	5.72 (5.98)	5.93 (6.34)	6.16 (6.79)
1	10.80 (9.25)	10.83 (9.38)	10.94 (9.82)	11.12 (10.63)	11.29 (11.59)
2	16.16 (12.55)	16.16 (12.71)	16.13 (13.24)	16.05 (14.24)	15.90 (15.41)
Weitzman Calibration					
0.5	10.54 (6.98)	11.81 (7.26)	16.50 (8.31)	18.24 (10.83)	21.06 (12.54)
1	18.73 (11.74)	19.69 (12.23)	24.21 (14.04)	42.86 (18.07)	72.44 (23.71)
2	24.58 (15.41)	25.08 (15.98)	26.81 (18.07)	51.14 (22.73)	89.93 (29.05)

Table 5: Sensitivity Analysis of SCC for Risk Aversion and EIS. The table shows SCC [\$/tCO₂] in 2015 for different values of γ and ψ . The numbers in front of the brackets are the results for the growth rate impact. The numbers in brackets are the results for the corresponding level impact.

effect on the growth rate is apparently too moderate, this property has a significant influence in the Weitzman calibration.

5 Robustness Checks

This section presents robustness checks for elasticity of intertemporal substitution, risk aversion, and diffusion parameters. We also compare our findings to the results in the DICE model.

5.1 Preference Parameters

Optimal Abatement and SCC We first consider the effect of varying the elasticity of intertemporal substitution, $\psi \in \{0.5, 1, 2\}$. Table 5 reports the results for the social cost of carbon and shows a strong dependence on EIS.²⁵ For a high level of EIS, society is willing to accept less smooth consumption streams. Consequently, it implements a more rigorous abatement policy raising SCC. The opposite is true for a low level of EIS. These results hold for both level and growth rate impact regardless of the calibration of the damages.

The effect of varying the degree of relative risk aversion depends on the damage specification and calibration. If damages are moderate for high temperatures (Nordhaus calibration), risk aversion is negligible in a model with a growth rate impact (G-N) and slightly more pronounced with a level impact (L-N). Nevertheless, the effects are relatively small. These results are in line with the findings of Ackerman et al. (2013) and Crost and Traeger (2014) that risk aversion has a much smaller effect than EIS on the optimal abatement decision and in turn on SCC.²⁶

²⁵This is in line with the findings of Cai et al. (2015), Crost and Traeger (2014), Jensen and Traeger (2014) and Bansal et al. (2014).

²⁶Crost and Traeger (2014) point out that most integrated assessment models are formulated for a CRRA decision maker with $\psi = 1/\gamma$. Since risk aversion plays an inferior role for the social cost of carbon and the optimal abatement policy, it is important to calibrate the entangled preference parameters to match EIS, rather than risk aversion. Especially for deterministic models, where risk aversion is in fact irrelevant, this might lead to significant changes in the optimal abatement policies.

ψ	χ		\mathcal{I}/C	
Nordhaus Calibration				
0.5	[75.4%, 81.9%]	(75.4%, 81.9%)	[18.1%, 24.2%]	(18.0%, 23.7%)
1	[75.0%, 75.0%]	(75.0%, 75.0%)	[24.0%, 25.0%]	(23.8%, 25.0%)
2	[72.2%, 74.7%]	(72.2%, 74.4%)	[24.5%, 27.7%]	(24.2%, 27.8%)
Weitzman Calibration				
0.5	[75.6%, 81.6%]	(75.6%, 82.0%)	[17.3%, 23.9%]	(18.0%, 24.2%)
1	[75.0%, 75.0%]	(75.0%, 75.0%)	[21.2%, 25.0%]	(22.3%, 25.0%)
2	[72.2%, 75.1%]	(72.2%, 74.6%)	[23.3%, 27.7%]	(24.1%, 27.8%)

Table 6: Sensitivity Analysis of Consumption and Investment for EIS. The table shows the range of optimal consumption and investment (as fraction of output) for different values of ψ . The numbers in box brackets are the results for the growth rate impact. The numbers in curved brackets are the results for the corresponding level impact.

However, if damages are potentially severe for high temperatures (Weitzman calibration), the results become sensitive to the choice of risk aversion for both damage specifications. Now, SCC and optimal abatement policy increase with risk aversion.²⁷

To summarize, abatement expenditures lead to steeper consumption streams (less consumption today, potentially more consumption in the future) and thus the EIS has a first-order effect. On the other hand, risk aversion or prudence become only relevant if the consequences of postponing abatement are severe and significantly state-dependent as in (L-W) and (G-W).

Optimal Consumption and Investment For unit EIS, the optimal consumption rates are constant. Lemma F.1 shows that for non-unit EIS the optimal consumption rate becomes state-dependent. Table 6 summarizes the effects of varying EIS on optimal consumption and investment, both expressed as a fraction of output.

We find that for $\psi > 1$, the optimal consumption rates are smaller than for unit EIS. Additionally to the more stringent abatement policy, society also installs more new capital via higher investment rates. Therefore, the gross growth rate of output is higher for $\psi > 1$. This confirms our intuition that with higher EIS, society accepts less smooth consumption streams, while the opposite is true for $\psi < 1$.²⁸

²⁷Our results also suggest that prudence, which is given by $\gamma(1 + \psi)$ (see Kimball and Weil (2009)), has a second-order effect as well. This is because prudence is affected similarly by risk aversion and EIS. If prudence had a significant effect on our results, then varying γ should also lead to significant changes, but this is only true when the consequences of postponing abatement are potentially disastrous.

²⁸Notice that for our benchmark choice of unit EIS, Section 3.4 calibrates $\vartheta = 0.372$ in order to match a consumption rate of 75%. If we choose ϑ to be 0.32(0.4) for an EIS of 0.5(2), then the consumption rate is in the range of 72%(74%) and 79%(76%), which is well in line with the historical range of 72% and 78%. More importantly, SCC for the different choices of ϑ are almost identical.

	σ_c	0	0.0081	0.0162	0.0243
Nordhaus Calibration		11.10 (10.61)	11.11 (10.62)	11.12 (10.63)	11.13 (10.64)
Weitzman Calibration		42.84 (18.05)	42.85 (18.06)	42.86 (18.07)	42.86 (18.09)
	σ_τ	0	0.05	0.1	0.15
Nordhaus Calibration		10.14 (7.44)	10.41 (8.26)	11.12 (10.63)	11.75 (13.81)
Weitzman Calibration		15.93 (8.67)	19.95 (10.88)	42.86 (18.07)	69.87 (33.03)

Table 7: SCC for Different Volatility Parameters. The table compares SCC [$\$/\text{tCO}_2$] for different volatility parameters for the four damage specifications. The results of the level specifications are in brackets.

5.2 Influence of Diffusive Shocks and Feedback Effects

Diffusive Shocks Table 7 shows how SCC in the year 2015 changes if the diffusion parameters of output and temperature, σ_c and σ_τ , are varied. It turns out that the volatility σ_c of economic shocks has a negligible effect on the current SCC. On the other hand, the effect of σ_τ is significant, since high variation in temperature amplifies the risk of ending up in a feedback loop during which temperature increases heavily. This is because the jump intensity increases in temperature. Therefore, society tries to avoid feedback loops by implementing a more rigorous abatement policy. Table 7 reports SCC for the four damage specifications. It can also be seen that SCC is more sensitive for the level impact.

Stochastic Feedback Effects We now analyze the impact of disregarding the stochastic feedback effects, i.e. $\pi_\tau(\tau) = \theta_\tau(\tau) = 0$. To obtain an expected equilibrium climate sensitivity of 3°C , we now choose $\eta_\tau = 4.33$. Notice that this specification can match the first two moments of ECS, but it cannot generate a fat-tailed climate sensitivity. For the Nordhaus damage specifications, SCC reduces from 11.12 (10.63) to 8.90 (4.81), where the number in brackets are the results for the level impact (L-N). Similar, for the Weitzman specifications, SCC in the year 2015 decreases from 42.86 (18.07) to 18.97 (5.34). We thus conclude that fat-tailed distributed climate dynamics induce a higher social cost of carbon and higher optimal abatement. The effect is more pronounced for level impacts where a climate feedback loop has potentially disastrous direct impacts on the economy.

5.3 Comparison with DICE

This subsection compares our benchmark results with those obtained in the DICE version of Nordhaus and Sztorc (2013). In particular, we compare the optimal social cost of carbon to Nordhaus' calculations. Nordhaus estimates the social cost of carbon in 2015 to be 19.6 dollars (expressed in 2005-dollars per ton of carbon dioxide). He uses a CRRA utility function with $\gamma = 1.45$ implying an EIS of $\psi = 1/\gamma$. By contrast, we use recursive preferences with $\gamma = 10$ and $\psi = 1$. The starting value of the social cost of carbon in our model is lower than estimated in

Abatement Policy		2015	2035	2055	2075	2095	2115	2150	2200
Optimal	GDP [trillion \$] (5% quantile)	75.8	124.1	195.4	284.7	386.8	501.4	724.7	1083.6
	GDP [trillion \$] (median)	75.8	139.3	230.4	348.7	490.2	652.7	988.0	1558.3
	GDP [trillion \$] (95% quantile)	75.8	156.5	272.1	428.0	620.7	852.6	1351.0	2244.8
	Temperature rise (5% quantile) [°C]	0.9	1.0	1.3	1.6	1.8	1.9	1.9	1.8
	Temperature rise (median) [°C]	0.9	1.3	1.7	2.1	2.4	2.5	2.7	3.1
	Temperature rise (95% quantile) [°C]	0.9	1.5	2.0	2.4	2.6	2.9	3.4	4.9
	Abatement Expenditures [trillion \$]	0.01	0.12	0.73	2.41	4.98	6.69	6.99	6.00
	Emission Control Rate	0.22	0.25	0.44	0.65	0.83	0.93	1	1
DICE	GDP [trillion \$] (5% quantile)	75.8	123.9	195.1	284.6	389.3	501.9	721.9	1061.4
	GDP [trillion \$] (median)	75.8	139.1	229.9	348.4	491.3	652.1	979.9	1536.2
	GDP [trillion \$] (95% quantile)	75.8	156.3	271.4	427.5	620.9	850.2	1335.9	2213.3
	Temperature rise (5% quantile) [°C]	0.9	1.0	1.2	1.4	1.4	1.2	0.9	0.6
	Temperature rise (median) [°C]	0.9	1.2	1.6	1.9	2.2	2.4	2.2	2.2
	Temperature rise (95% quantile) [°C]	0.9	1.5	2.0	2.6	3.1	3.5	4.4	6.4
	Abatement Expenditures [trillion \$]	0.05	0.24	0.82	2.20	4.91	8.45	7.74	6.13
	Emission Control Rate	0.20	0.32	0.54	0.62	0.81	1	1	1

Table 8: Optimal vs. DICE Abatement Policy for Level Impact. The table summarizes the simulation results obtained by running our model (L-N) with the optimal abatement policy and with the DICE abatement policy.

the latest version of DICE. In our model, however, society optimally anticipates environmental shocks and adjusts both the optimal abatement rate and the consumption rate. Along a path with high optimal abatement (as a response to high temperatures), the corresponding SCC values are significantly larger than the estimates in DICE. It is important to mention that DICE is formulated in a purely deterministic setting. In particular the temperature dynamics are calibrated to expected environmental outcomes, but do not take the uncertainty immanent in the climate system into account.

To analyze these points, we run our model with the optimal abatement policy obtained from DICE. Notice that following this policy is suboptimal in our model. The simulation results are summarized in Table 8. It turns out that the DICE abatement policy is more stringent than the median optimal policy. This leads to significant GDP losses, since the benefits of the DICE policy are lower than their abatement costs. Additionally, the DICE abatement policy is insensitive to unexpected variations in temperature, since it is determined in a deterministic model. By contrast, the optimal abatement policy reacts to high temperatures by tightening the abatement activities. This raises the social cost of carbon beyond the optimal value suggested by DICE. Conversely, along paths with low abatement, society raises consumption and SCC values are smaller. In contrast to the outcomes of following the (suboptimal) DICE policy, the variation of optimally controlled global temperatures and in turn the variation of climate damages is significantly smaller, while the variation of emissions is much higher.

6 Conclusion

This paper studies a flexible dynamic stochastic equilibrium model for optimal carbon abatement. All key variables such as carbon concentration, global temperature and world GDP are modeled as stochastic processes. Therefore, we can determine state-dependent optimal policies and provide model-based confidence bands for all our results. We perform a sophisticated calibration of all three model components (carbon concentration, global temperature, economy). In particular, we match the future distributions of transient climate response (TCR) and equilibrium climate sensitivity (ECS) as provided in the report of the IPCC (2014).

We study both a level and growth rate impact of temperatures on output and combine these two specifications with two alternative calibrations of the damage function. One calibration suggests moderate effects of climate change even for high temperature as in Nordhaus and Sator (2013), whereas the other calibration involves potentially severe damages as in Weitzman (2012). Therefore, we can compare four different scenarios: growth-rate impact and moderate damages (G-N), growth-rate impact and severe damages (G-W), level impact and moderate damages (L-N), level impact and severe damages (L-W)). We find that depending on the specification of the damage function the results can be very different. First, the social cost of carbon are similar for frameworks with level or growth rate impact if the potential damages of global warming are moderate. This changes significantly if damages can be severe. Then SCC is more than twice as large for a growth rate impact than for a level impact. The results are qualitatively similar for the optimal abatement policies.

We also document the effect of varying risk aversion and elasticity of intertemporal substitution on our results. If damages are moderate for high temperatures, risk aversion only matters when climate change has a level impact on output. Nevertheless, the effects are relatively small. By contrast, the elasticity of intertemporal substitution has a significant effect for both level and growth rate impact. If however damages are potentially severe for high temperatures, the results are sensitive to the choice of risk aversion for both impacts. Now, SCC and optimal abatement policy increase with risk aversion.

Finally, we find that in all scenarios the optimal abatement policies are state-dependent, but the strength of this dependence varies across scenarios. Given a Nordhaus damage calibration, the median results for optimal abatement policies and thus optimal emissions are similar, but the variations are higher for the level than for the growth-rate impact. In both cases, the optimal policies are less state-dependent than for the Weitzman damage calibration where the abatement policies are more rigorous.

References

- Ackerman, Frank, and Ramón Bueno, 2011, Use of McKinsey Abatement Cost Curves for Climate Economics Modeling, *Working Paper*, Stockholm Environment Institute.
- Ackerman, Frank, Elizabeth A. Stanton, and Ramón Bueno, 2011, CRED: A New Model of Climate and Development, *Ecological Economics*, 85, 166–176.
- Ackerman, Frank, Elizabeth A. Stanton, and Ramón Bueno, 2013, Epstein-Zin utility in DICE: Is risk aversion irrelevant to climate policy?, *Environmental and Resource Economics* 56, 73–84.
- Bansal, Ravi, Dana Kiku, and Marcelo Ochoa, 2014, Climate Change and Growth Risks, *Working Paper*, Duke University.
- Bansal, Ravi, and Marcelo Ochoa, 2011, Welfare Costs of Long-Run Temperature Shifts, *Working Paper*, NBER.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles., *The Journal of Finance* 1481–1509.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel, 2015, Global non-linear effect of temperature on economic production, *Nature* 527, 235–239.
- Cai, Yongyang, L. Judd Kenneth, and Thomas S. Lontzek, 2015, The Social Cost of Carbon with Economic and Climate Risks, *Working Paper*, NBER.
- Campbell, John Y., 1999, Asset prices, consumption, and the business cycle, in *J.B: Taylor, M. Woodford (eds.) Handbook of Macroeconomics*, volume 1 (Elsevier North-Holland).
- Crost, Benjamin, and Christian P. Traeger, 2014, Optimal CO₂ mitigation under damage risk valuation, *Nature Climate Change* 4, 631–636.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, 2009, Temperature and Income: Reconciling new Cross-sectional and Panel Estimates, *American Economic Review* 99, 198–204.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, 2012, Temperature Shocks and Economic Growth: Evidence from the Last Half Century, *American Economic Journal: Macroeconomics* 4, 66–95.
- Dietz, Simon, and Nicholas Stern, 2015, Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus’ Framework Supports Deep Cuts in Carbon Emissions, *Economic Theory* 125, 572–620.

- Duffie, Darrell, and Larry G. Epstein, 1992, Asset Pricing with Stochastic Differential Utility, *Review of Financial Studies* 5(3), 411–36.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinsky, 2014, Optimal Taxes on Fossil Fuel in General Equilibrium, *Econometrica* 82, 41–88.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–347.
- Hayashi, Fumio, 1982, Tobin’s marginal q and average q : A neoclassical interpretation, *Econometrica* 50, 213–224.
- Hope, Chris, 2006, The Marginal Impact of CO₂ from PAGE 2002: An Integrated Assessment Model Incorporating the IPCC’s Five Reasons for Concern, *The Integrated Assessment Journal* 6, 19–56.
- IPCC, 2014, *Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (Cambridge University Press).
- Jensen, S., and C.P. Traeger, 2014, Optimal Climate Change Mitigation under Long-Term Growth Uncertainty: Stochastic Integrated Assessment and Analytic Findings, *European Economic Review* 69(C), 104–125.
- Jermann, Urban J., 1998, Asset pricing in production economies, *Journal of Monetary Economics* 41, 257–275.
- Kelly, David L., and Charles D. Kolstad, 1999, Bayesian learning, growth, and pollution, *Journal of Economic Dynamics and Control* 23, 491–518.
- Kelly, David L., and Zhuo Tan, 2015, Learning and climate feedbacks: Optimal climate insurance and fat tails, *Journal of Environmental Economics and Management* 72, 98–122.
- Kimball, Miles, and Philippe Weil, 2009, Precautionary Saving and Consumption Smoothing across Time and Possibilities, *Journal of Money, Credit and Banking* 41, 245–284.
- McKinsey and Company, 2009, *Pathways to a Low-Carbon Economy: Version 2 of the Global Greenhouse Gas Abatement Cost Curve* (McKinsey and Company, London).
- McKinsey and Company, 2010, *Impact of the financial crisis on carbon economics: Version 2.1 of the Global Greenhouse Gas Abatement Cost Curve* (McKinsey and Company, London).
- Moore, Frances C., and Delavane B. Diaz, 2015, Temperature impacts on economic growth warrant stringent mitigation policy, *Nature Climate Change* 5, 127–131.
- Munk, Claus, and Carsten Sørensen, 2010, Dynamic Asset Allocation with Stochastic Income and Interest Rates, *Journal of Financial Economics* 96, 433–462.

- Nordhaus, William D., 1992, An Optimal Transition Path for Controlling Greenhouse Gases, *Science* 258, 1315–1319.
- Nordhaus, William D., 2008, *A Question of Balance: Weighing the Options on Global Warming Policies* (Yale University Press, New Haven).
- Nordhaus, William D., and Paul Sztorc, 2013, DICE 2013R: Introduction and Users Manual, *Technical Report*, Yale University.
- Pindyck, Robert S., 2011, Modeling the impact of warming in climate change economics, in *The Economics of Climate Change: Adaptations Past and Present*, 47–71 (NBER).
- Pindyck, Robert S., 2012, Uncertain Outcomes and Climate Change Policy, *Journal of Environmental Economics and Management* 63, 289–303.
- Pindyck, Robert S., 2013, Climate Change Policy: What Do the Models Tell Us?, *Journal of Economic Literature* 51, 860–872.
- Pindyck, Robert S., 2014, Risk and Return in the Design of Environmental Policy, *Journal of the Association of Environmental and Resource Economists* 1.
- Pindyck, Robert S., and Neng Wang, 2013, The Economic and Policy Consequences of Catastrophes, *American Economic Journal: Economic Policy* 5, 306–339.
- Stern, Nicholas, 2007, *The Economics of Climate Change: The Stern Review* (Cambridge University Press).
- Tol, Richard S.J., 2002, Estimates on the Damage Costs of Climate Change, Part I: Benchmark Estimates, *Environmental and Resource Economics* 21, 47–73.
- Traeger, Christian, 2014, A 4-stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change, *Environmental and Resource Economics* 59, 1–37.
- Vissing-Joergensen, Anette, 2002, Limited Asset Market Participation and the Elasticity of Intertemporal Substitution, *Journal of Political Economy* 110, 825–853.
- Vissing-Joergensen, Anette, and O. P. Attanasio, 2003, Stock-market Participation, Intertemporal Substitution, and Risk-aversion, *American Economic Review* 93, 383–391.
- Weitzman, Martin L., 2009, On Modeling and Interpreting the Economics of Catastrophic Climate Change, *The Review of Economics and Statistics* 91, 1–19.
- Weitzman, Martin L., 2012, GHG Targets as Insurance Against Catastrophic Climate Damages, *Journal of Public Economic Theory* 14 (2), 221–244.




Figure 6: Equilibrium Climate Sensitivity Graph (a) shows a histogram of the simulated ECS using the benchmark calibration. Graph (b) depicts a histogram of the ECS if we turn off the stochastic feedback effects. The histograms are based on a simulation of 1 million sample paths.

A Global Warming Process

The starting point for our climate model is the empirically well-documented logarithmic relationship between global warming and atmospheric carbon dioxide concentrations (see IPCC (2014)). A deterministic description of this relation is

$$T_t = \eta_\tau \log \left(\frac{Y_t^\Sigma}{Y^{\text{PI}}} \right). \quad (24)$$

Applying Ito's lemma to (24) and using (2) implies

$$dT_t = \frac{Y_t \eta_\tau}{Y_t^\Sigma} \left(\mu_y(t) - \alpha_t - \frac{1}{2} \frac{Y_t}{Y_t^\Sigma} \sigma_y^2 \right) dt + \frac{Y_t \eta_\tau}{Y_t^\Sigma} \sigma_y dW_t^y. \quad (25)$$

Notice that σ_y is empirically negligible. In Section 3.2, we use historical carbon dioxide data and estimate $\sigma_y = 0.0078$. This implies $|\frac{1}{2} \frac{Y_t}{Y_t^\Sigma} \sigma_y^2| \leq 3 \cdot 10^{-5}$, so that the term $-\frac{1}{2} \frac{Y_t}{Y_t^\Sigma} \sigma_y^2$ is close to zero. In the sequel, we thus drop this term.

Empirically, the relation between the temperature increase and carbon dioxide concentration is not deterministic (as assumed in (24)), but noisy. This calls for an additional modification wherefore we add two additional sources of randomness: First, we allow the temperature to be driven by a Brownian shock that is not necessarily perfectly correlated with W^y and that potentially induces more noise than the shock in (25), which is in line with empirical evidence. The latter means that we replace the diffusion parameter σ_y by σ_τ . Second, there is empirical evidence that the distribution of future temperature changes is right-skewed (see IPCC (2014)) as

a response to delayed climate feedback loops climate feedback loops. We thus add a self-exciting process that captures the idea of feedback loops. This is why we postulate the dynamics (8).

Figure 6 shows that the model is able to generate a right-skewed distribution of the equilibrium climate sensitivity. The histograms are based on simulations of one million sample paths. Graph (a) depicts a histogram of the equilibrium climate sensitivity using the benchmark calibration from Section 3. Graph (b) shows the corresponding histogram for a model without stochastic feedback effects. Our benchmark calibration yields a right-skewed distribution with a skewness of 0.92. By contrast, the skewness of the ECS in the model without stochastic feedback effects is approximately zero.

B Production Technology

Following Pindyck and Wang (2013), among others, we assume that the output C is generated by an AK -production technology

$$C = AK, \quad (26)$$

where K denotes the capital stock, which is the only factor of production. The parameter A denotes its productivity that is assumed to be constant. In this specification, K is the total stock of capital including physical, human, and firm-specific intangible capital. Following Nordhaus (2008), we assume that output can be used for investment \mathcal{I} , abatement expenditures \mathcal{A} and consumption \mathcal{C} , i.e. the restriction (12) holds

$$C = \mathcal{C} + \mathcal{I} + \mathcal{A}.$$

To express these values relative to output, we use the notation $\chi = \frac{\mathcal{C}}{C}$ for the consumption rate, $\nu = \frac{\mathcal{I}}{C}$ for the investment rate and $\varkappa = \frac{\mathcal{A}}{C}$ for the relative abatement expenditures. Therefore,

$$1 = \chi + \nu + \varkappa.$$

In a setting with *growth rate impact*, the dynamics of the capital stock are assumed to follow

$$\begin{aligned} dK_t &= (\mathcal{I}_t - \delta_k(t)K_t - \Phi(\nu_t, \varkappa_t, K_t) - \zeta_d T_t K_t) dt + \sigma_k K_t dW_t^k \\ &= K_t [(A\nu_t - \delta_k(t) - \phi(\nu_t + \varkappa_t) - \zeta_d T_t) dt + \sigma_k dW_t^k], \end{aligned}$$

where the costs $\Phi(\nu, \varkappa, K)$ of investment or abatement are homogenous of degree one:²⁹

$$\Phi(\nu, \varkappa, K) = \phi(\nu + \varkappa)K.$$

²⁹The homogeneity assumption is widely used in the literature. See, e.g., Pindyck and Wang (2013), Hayashi (1982) and Jermann (1998).

In the sequel, we assume that the cost function ϕ is quadratic, i.e. $\phi(1 - \chi) = 0.5\vartheta(1 - \chi)^2$. Capital grows due to investments \mathcal{I} . The combined effect that physical capital depreciates and human capital appreciates is captured by the (time-dependent) rate δ_k . Finally, capital is affected by economic shocks modeled by a Brownian motion W^k that can be correlated with W^τ and W^y . By (26), the output dynamics are then given by

$$dC_t = C_t[(A(1 - \chi_t) - \delta_k(t) - \phi(1 - \chi_t) - \zeta_d T_t)dt + \sigma_k dW_t^k - A\chi_t dt].$$

Comparing these dynamics with (13) and (17) yields (11) and

$$\kappa = A\chi.$$

Therefore, (14) follows.

For the specification with a *level impact*, we can repeat the same steps as above, but replace all economic variables by the corresponding shadow variables (e.g. C by \widehat{C}). We assume $\mathcal{C}_t = \widehat{\mathcal{C}}_t D_t \mathcal{K}_t$, i.e. $\widehat{\chi} = \chi$. This assumption is satisfied for unit EIS and is approximately true in the general case. It thus follows $\mathcal{I}_t + \mathcal{A}_t = \widehat{\mathcal{I}}_t D_t \mathcal{K}_t$, which implies (14).³⁰

C Calibration of Natural Sinks

We now provide the details on how to calibrate the natural carbon dioxide sinks. We use annual RCP 8.5 emission data denoted by $\{\widehat{E}_n\}_{n=1}^N$ at times $\{t_n\}_{n=1}^N$. This data represents our BAU-emissions. Additionally, we use RCP 8.5 concentration data $\{Y_n^{\text{BAU}}\}_{n=1}^N$ at the same time points. In the first step, we solve (5) for δ_y and use that data to calculate a set of model-implied carbon dioxide decay rates via

$$\delta_y(X_n^{\text{BAU}}) = \frac{\xi \widehat{E}_n - Y_n^{\text{BAU}} \mu_y(t_n)}{Y_n^{\text{BAU}}}, \quad n = 1, \dots, N.$$

Since the RCP 8.5 emission and concentration data is only available for steps of five years, we interpolate the decay rate and the concentration between the data points via cubic splines. Then we simulate the process X_t^{BAU} applying the Euler method to (4) with one time step per year:

$$X_{t+1}^{\text{BAU}} = X_t^{\text{BAU}} + \delta_y(X_t^{\text{BAU}}) Y_t^{\text{BAU}}.$$

Here, we choose without loss of generality $X_0^{\text{BAU}} = 0$ as the starting point of the process. We obtain pairs of annual data points $\{X_t^{\text{BAU}}, \delta_y(X_t^{\text{BAU}})\}$. It turns out that these decay rates can

³⁰Our simulations show that the error is negligible for $\psi \neq 1$. A detailed derivation is available from the authors upon request.

be fitted well using the following functional form:

$$\delta_y(x) = a_\delta e^{-\left(\frac{x-b_\delta}{c_\delta}\right)^2},$$

where $a_\delta = 0.017$, $b_\delta = 11.64$, $c_\delta = 279.7$. This parametrization achieves an almost perfect fit to the data ($R^2 > 99\%$), and Graph (b) in Figure 2 shows that the model excellently replicates the RCP 8.5 emission data.

D Calibration of the Economic Model

GDP Growth We calibrate the growth rate (11) such that our economic model closely matches the evolution of GDP growth in the latest version of the DICE model.³¹ Notice that (11) is the growth rate before abatement and climate damage. Therefore, we simulate the output in the DICE model disregarding abatement and climate damage.³² This yields data points $(t_1, \hat{C}_1), \dots, (t_n, \hat{C}_n)$, which are used to extract the corresponding future GDP growth rates of DICE. It turns out that these growth rates (before abatement and damages) can be fitted well using the following functional form:

$$g^{\text{DICE}}(t) = g_0 + g_1 e^{-\delta_g t} \quad (27)$$

where $g_0 = 0.005$, $g_1 = 0.029$, $\delta_g = 0.011$, i.e. the growth rate declines at a rate of 1.1% to its long-term steady-state level of 0.5%. Figure 7 depicts the simulated data points of the DICE model (crosses) and the fitted values (solid line).

In a second step, we equate (11) with (27) and determine the unknowns in (11). Following Pindyck and Wang (2013), we fix the productivity at $A = 0.113$. To separately identify $\delta_k(t)$ and ϑ , we must make an assumption about the consumption rate χ , which is an endogenous variable. The optimal χ in DICE is pretty stable over time and close to 75%. This is also in line with historical data. Since unit EIS is our benchmark choice, the optimal consumption rate is a constant. It can be determined from the following quadratic equation³³

$$\delta = \chi^*(A - \vartheta(1 - \chi^*)).$$

Therefore, assuming that the optimal consumption rate of the society is $\chi^* = 0.75$ we obtain

³¹Notice that the GDP growth in DICE is calibrated to historical data along with projected future growth rates of population, production and technology. In order to simplify the comparison with DICE, we have decided to match the future evolution of world GDP. The DICE model is however deterministic. So matching means that we match the average evolution in our model to the DICE dynamics.

³²In the terminology of our paper, this is the shadow GDP of DICE.

³³See the first-order condition (34) in Appendix F. Notice that for unit EIS this first-order condition does not depend on abatement or climate change. In the robustness section, we study cases where EIS is not one.




Figure 7: Calibration of the GDP Process. The figure depicts GDP forecasts simulated by the DICE model (crosses) and our median path based on simulations with the estimated parameters (solid line).

$\vartheta = 0.372$ leading to

$$\delta_k(t) = 0.0116 - 0.029e^{-0.011t}.$$

To determine the volatility of the GDP process, we cannot use the DICE model which is deterministic. We thus use historical data from the website of the International Monetary Fund starting in 1960 to estimate this parameter.³⁴ Calculating the standard deviation of the log returns yields a volatility of $\sigma_c = 0.0161$. Furthermore we obtain the following correlations with global temperature and carbon dioxide $\rho_{cT} = 0$ and $\rho_{cY} = 0.29$.

Abatement Costs Following Nordhaus (2008) and Nordhaus and Sztorc (2013) we assume that abatement expenditures are proportional to the current level of GDP:

$$\mathcal{A}_t^\alpha = C_t \theta_1(t) (\varepsilon_t^\alpha)^{\theta_2} \quad (28)$$

for a constant $\theta_2 > 1$ and a time dependent function θ_1 that gradually decreases over time. The decreasing trend in abatement expenditures reflects the widening menu of sustainable technological alternatives implying that abatement becomes cheaper over time. We adopt the parametrization from the DICE model and use $\theta_1(t) = 0.05506 \exp(-0.0148t) + 0.00043$ and $\theta_2 = 2.8$. Combining (14) and (28) yields the following instantaneous growth effect of abatement

$$\kappa(t, \varepsilon_t^\alpha) = A \theta_1(t) (\varepsilon_t^\alpha)^{\theta_2}. \quad (29)$$

³⁴Available at: <https://www.imf.org/external/data.htm>

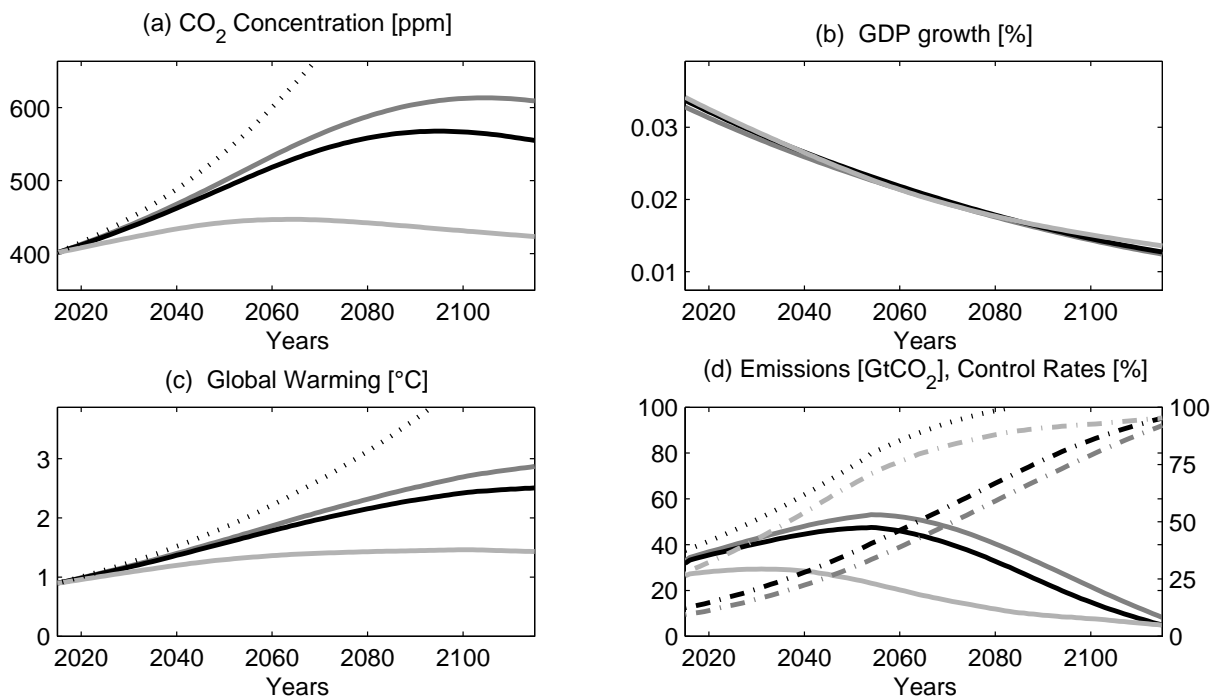


Figure 8: Sensitivity Analysis for the Preferences. The graphs show the median paths of the key variables for different preference specifications. Median optimal paths are depicted by solid lines and median BAU paths by dotted lines. The benchmark scenario is depicted by black lines. Grey lines show the DICE preference structure and light lines represent Stern discounting. Graph (a) shows the carbon dioxide concentrations in the atmosphere, (b) median GDP growth rates, (c) median changes in global temperature, (d) carbon dioxide emissions and the optimal emission control rates (dash-dotted lines).

E Further Robustness Checks

E.1 Comparison with Standard Preference Choices

We now compare our benchmark preference structure with two specifications that are standard in the literature. First, we consider a time-additive CRRA utility function with a risk aversion parameter of $\gamma = 1.45$ and time-preference rate of $\delta = 1.5\%$. This utility specification is used as benchmark specification in DICE. Earlier versions and other models use similar CRRA specifications.³⁵ Second, we simulate our model using a log-utility function ($\psi = \gamma = 1$) with a very low discount rate ($\delta = 0.1\%$). Pindyck (2013), among others, argues that optimal abatement policies crucially depend on the time-preference rate. In general, there is a lot of debate about this parameter in the IAM literature. This is because time-preferences put implicitly weights on the current and future generations: A higher value puts more weight on the current generation, whereas a lower value shifts some of this weight to future generations. A tension arises since the current generation is not as severely affected by the climate change as the future generations,

³⁵See, e.g., Nordhaus (2008), Pindyck (2012), Ackerman et al. (2011).

Preferences	2015	2035	2055	2075	2095	2115	2150	2200
Benchmark	11.12	21.75	50.67	102.52	171.21	225.10	254.12	353.25
DICE	7.67	14.03	35.15	77.01	137.83	201.57	243.49	324.95
Stern	38.11	77.30	139.02	185.47	208.60	223.44	250.53	377.84

Table 9: SCC for Different Preference Specifications. The table compares SCC [\$/tCO₂] for different preference specifications. The results are generated using the (G-N) damage specification. DICE preferences are $\gamma = 1.45$, $\psi = 1/\gamma$, $\delta = 1.5\%$. Stern refers to $\gamma = 1$, $\psi = 1$, $\delta = 0.1\%$.

but must today decide upon an optimal abatement policy and pay for it. Of course, more stringent actions reduce current consumption, but have far reaching consequences for future generations who might benefit the most. We refer to the very low discount rate of $\delta = 0.1\%$ in combination with log-utility as *Stern discounting* since Stern (2007) suggests this preference structure. Intuitively, with such a low rate of time-preference, real interest rates are lower and, in turn, the social cost of carbon is higher. This also implies that society implements a very stringent abatement policy.

Table 9 and Figure 8 summarizes our findings on how the preference structure affects our results. With the standard DICE preference structure, both risk aversion and EIS are lower which leads to a moderate abatement policy. The resulting evolution of median global warming is in line with the results presented in Nordhaus and Sztorc (2013) although SCC is significantly smaller. *Stern discounting* yields a very stringent abatement policy and a high social cost of carbon. Following this policy reduces carbon dioxide emissions so that the median temperature increase peaks by the end of this century at 1.5°C. Additionally, climate variability is significantly dampened (not shown in the figure).

E.2 Alternative Abatement Costs

We now study the effects of using an alternative cost function κ . Instead of the benchmark specification, we derive an abatement cost function using the prognosis for the marginal greenhouse gas abatement costs for the year 2030 provided by McKinsey and Company (2009, 2010).

Calibration The calibration is based on a prognosis for the marginal greenhouse gas abatement costs for the year 2030 provided by McKinsey and Company (2009, 2010). For that year, they estimate that under BAU the total emissions of greenhouse gases would reach 66GtCO₂e and analyze the expected abatement expenditures. Under rather optimistic assumptions, they report an abatement potential of 38GtCO₂e at a total cost of 150 billion euros. McKinsey supposes that for 11GtCO₂e of abatement the net costs are negative because savings from implementing energy-efficient measures – compared to the BAU scenario – exceed the initial investment costs. To avoid issues arising from negative abatement costs, we follow Ackerman and Bueno (2011) and disregard the negative part of the marginal costs. Therefore, our calibration




Figure 9: Calibration of the Abatement Costs. The figure depicts the marginal abatement cost (MAC) for the reference year 2030 (solid line). The prices of the y -axis are in 2005 euros. GtCO₂e stands for gigatons of carbon dioxide equivalents. The MAC function is calibrated such that it fits the positive part of the McKinsey data (crosses).

is more conservative than the McKinsey prognosis.

In a first step, we fit the McKinsey data using the following functional form for the marginal abatement cost function:

$$\text{MAC}(q) = \frac{c_1 q}{c_2 + c_3 q + c_4 q^2}.$$

The variable q is the absolute quantity of greenhouse gas abatement (measured in GtCO₂) compared to the business-as-usual scenario, i.e. the difference between BAU-emissions and controlled emissions, $q = \mathbf{e}^{\text{BAU}} - \mathbf{e}^\alpha$. As can be seen in Figure 9, our estimates of c_i fit the positive part of the marginal abatement costs well ($R^2 > 0.96$). The coefficients are $c_1 = 0.00039$, $c_2 = 0.0016$, $c_3 = -3.25 \cdot 10^{-5}$, $c_4 = -7.27 \cdot 10^{-8}$. Then, we transform the marginal costs MAC into (absolute) expenditures, which in our paper are denoted by \mathcal{A} . We thus compute the anti-derivative $\mathcal{A}(q)$ of the marginal costs with respect to q and evaluate \mathcal{A} at the available data points q_1, \dots, q_n . This yields values $\mathcal{A}_1, \dots, \mathcal{A}_n$.

The resulting data points $(q_1, \mathcal{A}_1), \dots, (q_n, \mathcal{A}_n)$ can now be used to determine the cost function $\kappa(t, \varepsilon^\alpha)$ for the year 2030. Notice that the McKinsey data maps absolute quantities of abatement q into marginal expenditures, whereas our cost function maps emission control rates into reductions of economic growth. Therefore, we transform absolute quantities of greenhouse gas abatement q into emission control rates using $\varepsilon_i = q_i / \mathbf{e}^{\text{BAU}}$, $i = 1, \dots, n$, and absolute abatement expenditures \mathcal{A} in relative expenditures by $\varkappa_i = \mathcal{A}_i / \mathbb{E}[C]$, $i = 1, \dots, n$, where $\mathbb{E}[C]$ denotes the expected GDP in 2030. We assume the functional form (28) still to hold. We calibrate the parameters such that (28) is close to the data points $(\varepsilon_1, \varkappa_1), \dots, (\varepsilon_n, \varkappa_n)$. As a result of the




Figure 10: Sensitivity Analysis for the Cost Function. The graphs show the median paths of the key variables for different specifications of the cost function. Median optimal paths are depicted by solid lines and median BAU paths by dotted lines. The benchmark scenario is depicted by black lines. Grey lines show the results using McKinsey abatement costs. Graph (a) shows the carbon dioxide concentrations, (b) median GDP growth rates, (c) median changes in global temperature, (d) carbon dioxide emissions and the optimal emission control rates (dash-dotted lines).

calibration we obtain $\theta_1 = 0.035$, $\theta_2 = 3.186$ ($R^2 > 99\%$) for the year 2030. We take the rate at which abatement becomes cheaper over time from DICE, i.e. the relative expenditures for complete abatement ($\varepsilon = 1$) decline at rate of 1.48% to its long-term level of 0.043%. As a result, we obtain $\theta_1(t) = 0.0443 \exp(-0.0148t) + 0.00043$. Notice that the calibration based on the McKinsey prognosis makes abatement slightly cheaper than in DICE since θ_1 is smaller.

Results for (G-N) Figure 10 depicts the median results for both cost specifications and (G-N).³⁶ It turns out that the results are similar. Implementing the McKinsey specification, slightly raises the optimal abatement policy leading to lower carbon dioxide concentrations and temperatures compared to the benchmark case. Therefore, net GDP growth is slightly higher and SCC in the year 2015 is reduced from \$11.12 to \$10.05. Notice that, by the end of this century, the optimal abatement activity becomes higher for the DICE cost function. This is because for high emission control rates the marginal costs of the McKinsey calibration are higher than for DICE since θ_2 is higher (3.186 instead of 2.8). In turn, the marginal benefits from abatement are lower when control rates are high.

³⁶The results for (L-N) are similar and available from the authors upon request.

F Solution Method

The optimization problem (20) cannot be solved explicitly.³⁷ This appendix summarizes how the problem can be solved numerically.

F.1 Hamilton-Jacobi-Bellman Equation

In case of a growth rate impact, the Hamilton-Jacobi-Bellman (HJB) equation reads

$$\begin{aligned}
0 = \sup_{\alpha, \chi} & \left\{ J_t + c \left(g(t, \chi) - \zeta_d \tau^k - \kappa(t, \varepsilon^\alpha) \right) J_c + \frac{1}{2} c^2 \sigma_c^2 J_{cc} + y(\mu_y(t) - \alpha) J_y + \frac{1}{2} y^2 \sigma_y^2 J_{yy} \right. \\
& - \delta_y(x) y J_x + \frac{y \eta_\tau}{y + Y^{\text{PI}}} \left(\mu_y(t) - \alpha \right) J_\tau + \frac{1}{2} \left(\frac{\sigma_\tau y}{y + Y^{\text{PI}}} \right)^2 J_{\tau\tau} + \frac{y^2}{y + Y^{\text{PI}}} \rho_{y\tau} \sigma_y \sigma_\tau J_{y\tau} \\
& \left. + c y \rho_{cy} \sigma_y \sigma_c J_{cy} + c \sigma_c \rho_{c\tau} \frac{\sigma_\tau y}{y + Y^{\text{PI}}} J_{c\tau} + \pi_\tau(\tau) \left[J(t, c, y, \tau + \theta_\tau) - J \right] + f(\chi c, J) \right\}. \quad (30)
\end{aligned}$$

Subscripts denote partial derivatives (e.g. $J_t = \partial J / \partial t$). The corresponding HJB equation for a level impact can be found in Section F.3. First, we establish the following separation result:

Lemma F.1. *The indirect utility function of the optimization problem (20) has the form*

$$J(t, c, x, y, \tau) = \frac{1}{1 - \gamma} c^{1 - \gamma} F(t, x, y, \tau), \quad (31)$$

where F solves the simplified HJB equation

$$\begin{aligned}
0 = \sup_{\alpha, \chi} & \left\{ F_t + y \left[\mu_y(t) - \alpha + (1 - \gamma) \sigma_c \sigma_y \rho_{cy} \right] F_y + \frac{1}{2} y^2 \sigma_y^2 F_{yy} + \delta_y(x) y F_x \right. \\
& + \frac{y \eta_\tau}{y + Y^{\text{PI}}} \left[\mu_y(t) - \alpha + (1 - \gamma) \rho_{c\tau} \sigma_c \sigma_\tau \right] F_\tau + \frac{1}{2} \left(\frac{y \sigma_\tau}{y + Y^{\text{PI}}} \right)^2 F_{\tau\tau} \\
& + \frac{y^2}{y + Y^{\text{PI}}} \rho_{y\tau} \sigma_y \sigma_\tau F_{y\tau} + (1 - \gamma) \left[g(t, \chi) - \zeta_d \tau - \kappa(t, \varepsilon^\alpha) - \frac{1}{2} \gamma \sigma_c^2 - \frac{\delta}{1 - 1/\psi} \right] F \\
& \left. + \pi_\tau(\tau) \left[F(t, y, \tau + \theta_\tau(\tau)) - F \right] + \delta \theta \chi^{1 - 1/\psi} F^{1 - 1/\theta} \right\}, \quad (32)
\end{aligned}$$

The optimal abatement strategy is given by

$$\alpha_t^* = \kappa_\alpha(t, y, \cdot)^{-1} \left(\frac{y F_y + \frac{y \eta_\tau}{y + Y^{\text{PI}}} F_\tau}{(\gamma - 1) F} \right), \quad (33)$$

³⁷Notice that closed-form solutions are only available in rare special cases. A prominent example is the combination of log-utility, a Cobb-Douglas production technology and some further debatable assumptions as in Golosov et al. (2014). As discussed in Section 5, log-utility is too restrictive for studying the effects of preference parameters.

and the optimal consumption rate satisfies

$$\delta(\chi_t^*)^{-1/\psi} F^{-1/\theta} = -\frac{\partial}{\partial \chi} g(t, \chi_t^*). \quad (34)$$

Proof. Substituting the conjecture into the HJB equation yields the simplified HJB equation (32). The representations of the optimal controls are then obtained by calculating the first-order conditions. \square

The HJB equation cannot be simplified further. Therefore, we have to determine F by solving equation (32) numerically. First, we consider a simplified problem where the capacity of natural sinks is assumed to be unconstrained, i.e. the decay rate of carbon dioxide is assumed to be constant at $\bar{\delta}_y = \delta_y(0)$. This provisional assumption makes the state variable X redundant and significantly simplifies the solution algorithm. In a second step, we address the general case.

F.2 Numerical Solution Approach

Basic Idea We use a grid based solution approach to solve the non-linear PDE. We discretize the (t, y, τ) -space using an equally spaced lattice. Its grid points are defined by

$$\{(t_n, y_i, \tau_j) \mid n = 0, \dots, M_t, i = 0, \dots, M_y, j = 0, \dots, M_\tau\},$$

where $t_n = n\Delta_t$, $y_i = i\Delta_y$, and $\tau_j = j\Delta_\tau$ for some fixed grid size parameters Δ_t , Δ_y , and Δ_τ that denote the distances between two grid points. The numerical results are based on a choice of $M_y = 500$, $M_\tau = 1000$ and 1 time step per year. Our results hardly change if we use a finer grid or more time steps per year. The parameters M_τ and M_y are chosen sufficiently large such that it is very unlikely that these boundaries are reached within the given time horizon. In the sequel, $F_{n,i,j}$ denotes the approximated indirect utility function at the grid point (t_n, y_i, τ_j) and $\alpha_{n,i,j}$ refers to the corresponding optimal abatement policy. We apply an implicit finite difference scheme.

Terminal Condition Since the optimization problem (20) has an infinite time horizon, we must transform it into a problem with a finite horizon. Therefore, we approximate the indirect utility function at some point $t_{\max} = M_t\Delta_t$ in the distant future – we choose the year 2500 – by the solution of a similar problem where the world is in a steady state: We assume that from time t_{\max} onwards the emission control rate is one, i.e. anthropological carbon dioxide emissions are zero. To approximate the indirect utility function at the grid point (t_{\max}, y_i, τ_j) , we simulate sample paths for further 500 years and determine the utility index (18).

Finite Differences Approach In this paragraph, we describe the numerical solution approach in more detail. We adapt the numerical solution approach used by Munk and Sørensen (2010).

The numerical procedure works as follows. At any point in time, we make a conjecture for the optimal abatement policy $\alpha_{n,i,j}^*$. A good guess is the value at the previous grid point since the abatement strategy varies only slightly over a small time interval, i.e. we set $\alpha_{n,i,j}^* = \alpha_{n+1,i,j}$. Substituting this guess into the HJB equation yields a semi-linear PDE:

$$0 = F_t + K_1 F^{1-\frac{1}{\theta}} + K_2 F + K_3 F_y + K_4 F_{yy} + K_5 F_\tau + K_6 F_{\tau\tau} + K_7 F_{\tau y} + \pi_\tau F(t, c, \tau + \theta_\tau)$$

with state dependent coefficients $K_i = K_i(t, y, \tau)$. Due to the implicit approach, we approximate the time derivative by forward finite differences. In the approximation, we use the so-called 'up-wind' scheme that stabilizes the finite differences approach. Therefore, the relevant finite differences at the grid point (n, i, j) are given by

$$\begin{aligned} D_y^+ F_{n,i,j} &= \frac{F_{n,i+1,j} - F_{n,i,j}}{\Delta_c}, & D_y^- F_{n,i,j} &= \frac{F_{n,i,j} - F_{n,i-1,j}}{\Delta_c}, \\ D_\tau^+ F_{n,i,j} &= \frac{F_{n,i,j+1} - F_{n,i,j}}{\Delta_\tau}, & D_\tau^- F_{n,i,j} &= \frac{F_{n,i,j} - F_{n,i,j-1}}{\Delta_\tau}, \\ D_{yy}^2 F_{n,i,j} &= \frac{F_{n,i+1,j} - 2F_{n,i,j} + F_{n,i-1,j}}{\Delta_c^2}, & D_{\tau\tau}^2 F_{n,i,j} &= \frac{F_{n,i,j+1} - 2F_{n,i,j} + F_{n,i,j-1}}{\Delta_\tau^2}, \\ D_t^+ F_{n,i,j} &= \frac{F_{n+1,i,j} - F_{n,i,j}}{\Delta_t}, & D_{\tau y}^2 F_{n,i,j} &= \frac{F_{n,i,j+1} - F_{n,i-1,j+1} - F_{n,i+1,j-1} + F_{n,i-1,j-1}}{4\Delta_\tau\Delta_y}. \end{aligned}$$

We approximate the jump terms via linear interpolation between the closest grid points:

$$F(t, y, \tau + \theta_\tau) = k_{\tau 1} F_{n,i,j+\hat{\theta}_{\tau 1}} + k_{\tau 2} F_{n,i,j+\hat{\theta}_{\tau 2}},$$

where $\hat{\theta}_{\tau 1}$ and $\hat{\theta}_{\tau 2}$ denote the closest grid points of $\tau + \theta_\tau$. The variables k_τ denote the weights resulting from linear interpolation. Substituting these expressions into the PDE above yields the following semi-linear equation for the grid point (t_n, y_i, τ_j)

$$\begin{aligned} F_{n+1,i,j} \frac{1}{\Delta_t} &= F_{n,i,j} \left[-K_2 + \frac{1}{\Delta_t} + \text{abs} \left(\frac{K_3}{\Delta_y} \right) + \text{abs} \left(\frac{K_5}{\Delta_\tau} \right) + 2 \frac{K_4}{\Delta_y^2} + 2 \frac{K_6}{\Delta_\tau^2} \right] \\ &+ F_{n,i-1,j} \left[\frac{K_3^-}{\Delta_y} - \frac{K_4}{\Delta_y^2} \right] + F_{n,i+1,j} \left[-\frac{K_3^+}{\Delta_y} - \frac{K_4}{\Delta_y^2} \right] \\ &+ F_{n,i,j-1} \left[\frac{K_5^-}{\Delta_\tau} - \frac{K_6}{\Delta_\tau^2} \right] + F_{n,i,j+1} \left[-\frac{K_5^+}{\Delta_\tau} - \frac{K_6}{\Delta_\tau^2} \right] \\ &+ F_{n,i-1,j+1} \frac{K_7}{4\Delta_\tau\Delta_y} + F_{n,i+1,j-1} \frac{K_7}{4\Delta_\tau\Delta_y} - F_{n,i+1,j+1} \frac{K_7}{4\Delta_\tau\Delta_y} - F_{n,i-1,j-1} \frac{K_7}{4\Delta_\tau\Delta_y} \\ &+ \pi_\tau (k_{\tau 1} F_{n,i,j+\hat{\theta}_{\tau 1}} + k_{\tau 2} F_{n,i,j+\hat{\theta}_{\tau 2}}) - K_1 F_{n,i,j}^{1-\frac{1}{\theta}}. \end{aligned}$$

Therefore, for a fixed point in time each grid point is determined by a non-linear equation. This results in a non-linear system of $(M_y + 1)(M_\tau + 1)$ equations that can be solved for the vector

$$F_n = (F_{n,1,1}, \dots, F_{n,1,M_\tau}, F_{n,2,1}, \dots, F_{n,2,M_\tau}, \dots, F_{n,M_y,1}, \dots, F_{n,M_y,M_\tau}).$$

Notice that in case of CRRA utility the system becomes linear. Using this solution we update our conjecture for the optimal abatement policy at the current point in the time dimension. We apply the first-order condition (33) and finite difference approximations of the corresponding derivatives. In the interior of the grid, we use centered finite differences. At the boundaries, we apply forward or backward differences. For instance, for $(i, j) \in \{2, \dots, M_y - 1\} \times \{2, \dots, M_\tau - 1\}$, we compute the new guess as

$$\alpha_{n,i,j}^* = \kappa_\alpha(t_n, y_i, \cdot)^{-1} \left(\frac{(y_i + Y^{\text{PI}})\Delta_\tau y_i (F_{n,i+1,j} - F_{n,i-1,j}) + \Delta_y y_i (F_{n,i,j+1} - F_{n,i,j-1})}{\Delta_y \Delta_\tau (y_i + Y^{\text{PI}})(\gamma - 1)F_{n,i,j}} \right).$$

Similarly, we compute the social cost of carbon for fixed GDP C in the grid point $(i, j) \in \{2, \dots, M_y - 1\} \times \{2, \dots, M_\tau - 1\}$ as

$$\text{SCC}_{n,i,j} = \frac{C}{1 - \gamma} \frac{F_{n,i+1,j} - F_{n,i-1,j}}{\Delta_y F_{n,i,j}} \frac{\xi_\epsilon}{\mu_y + \delta_y - \alpha_{n,i,j}^*}.$$

With this new guess for the optimal policy we perform a new iterative step. We continue the iteration until there is no significant change of the result. Then the algorithm continues with the previous point t_{n-1} in the time directions until we reach the end of the grid.

Implementation of State-Dependent Sinks The solution procedure described so far does not deal with state dependent sinks. Since in general the constraint (7) involves X , we first solve for the optimal abatement policy if the weaker constraint $\alpha_t \leq \mu_y(t) + \bar{\delta}_y$ is imposed. The corresponding abatement decision is then given by

$$\bar{\alpha}_t = \min \left[\mu_y(t) + \bar{\delta}_y, \kappa_\alpha(t, y, \cdot)^{-1} \left(\frac{yF_y + \frac{y\eta}{y+Y^{\text{PI}}}F_\tau}{(1 - \gamma)F} \right) \right].$$

Since the modified constraint is always weaker, we obtain an upper bound $\bar{J}(t, c, x, y, \tau) \geq J(t, c, x, y, \tau)$ for the indirect utility function of the true model where (7) is imposed. Of course, $\bar{\alpha}_t$ is not feasible in the true model. To obtain a feasible strategy, we thus define

$$\underline{\alpha}_t = \min [\mu_y(t) + \delta_y(X_t), \bar{\alpha}_t],$$

where we cut off $\bar{\alpha}_t$ if it violates (7). Notice that the strategy $\underline{\alpha}_t$ is suboptimal. Since we have the upper bound \bar{J} , we can compute an upper bound on the loss that occurs if we implement $\underline{\alpha}_t$ instead of the (unknown) optimal strategy. If $\underline{J}(t, c, x, y, \tau)$ denotes the indirect utility associated

with $\underline{\alpha}_t$, the upper bound on the welfare loss is given by

$$\underline{J}(t, c, x, y, \tau) = \bar{J}(t, c(1 - L), x, y, \tau).$$

It turns out that this upper bound for the welfare loss is significantly below 0.1% and thus the strategy $\underline{\alpha}_t$ is close to optimal.

Comparison with Value Function Iteration Most IAMs are formulated in discrete time. The corresponding Bellman equation is usually solved by dynamic programming with value function iteration (see, e.g., Crost and Traeger (2014), Traeger (2014), Cai et al. (2015)). The main idea is as follows: One first discretizes the state space and chooses an appropriate functional form to approximate the value function in those nodes. A typical choice are multivariate orthogonal polynomials (e.g. Chebychev polynomials). Starting from the terminal date, one iterates backwards through time. For every t_n one pointwise maximizes the right-hand side of the Bellman equation and determines the approximate value function in every node. Then the algorithm goes step-by-step back in time until the end of the time grid is reached.

Since our model is formulated in continuous time, we derive and solve the corresponding Hamilton-Jacobi-Bellman equation. This partial differential equation can be solved numerically by a grid-based finite-differences approach as described above. Our method is thus the continuous-time analogue to discrete-time value function iteration. Notice that it is not necessary to make any assumptions on the functional form of the value function.

F.3 Level Impact

For a level impact, we use the notation $\tilde{C}_t = \hat{C}_t \mathcal{K}_t^\alpha$. Then, the HJB equation reads

$$\begin{aligned} 0 = \sup_{\alpha, \chi} \Big\{ & J_t + \tilde{c} (g(t, \chi) - \kappa(t, \varepsilon^\alpha)) J_{\tilde{c}} + \frac{1}{2} \tilde{c}^2 \sigma_c^2 J_{\tilde{c}\tilde{c}} + y(\mu_y(t) - \alpha) J_y + \frac{1}{2} y^2 \sigma_y^2 J_{yy} - \delta_y(x) y J_x \\ & + \tilde{c} y \rho_{cy} \sigma_y \sigma_c J_{\tilde{c}y} + \frac{y \eta_\tau}{y + Y^{\text{PI}}} (\mu_y(t) - \alpha) J_\tau + \frac{1}{2} \left(\frac{\sigma_\tau y}{y + Y^{\text{PI}}} \right)^2 J_{\tau\tau} \\ & + \frac{y^2}{y + Y^{\text{PI}}} \rho_{y\tau} \sigma_y \sigma_\tau J_{y\tau} + \tilde{c} \sigma_c \rho_{\tilde{c}\tau} \frac{\sigma_\tau y}{y + Y^{\text{PI}}} J_{\tilde{c}\tau} + \pi_\tau(\tau) [J(t, \tilde{c}, y, \tau + \theta_\tau) - J] \\ & + f(\tilde{c} D(\tau) \chi, J) \Big\}. \end{aligned}$$

Lemma F.1 is then modified as follows:

Lemma F.2. *The indirect utility function of the optimization problem has the form*

$$J(t, \tilde{c}, x, y, \tau) = \frac{1}{1 - \gamma} \tilde{c}^{1 - \gamma} F(t, x, y, \tau),$$

where F solves the simplified HJB equation

$$\begin{aligned}
0 = \sup_{\alpha, \chi} & \left\{ F_t + y \left[\mu_y(t) - \alpha + (1 - \gamma) \sigma_c \sigma_y \rho_{cy} \right] F_y + \frac{1}{2} y^2 \sigma_y^2 F_{yy} + \delta_y(x) y F_x \right. \\
& + \frac{y \eta_\tau}{y + YPI} \left[\mu_y(t) - \alpha + (1 - \gamma) \rho_{c\tau} \sigma_c \sigma_\tau \right] F_\tau + \frac{1}{2} \left(\frac{y \sigma_\tau}{y + YPI} \right)^2 F_{\tau\tau} \\
& + \frac{y^2}{y + YPI} \rho_{y\tau} \sigma_y \sigma_\tau F_{y\tau} + (1 - \gamma) \left[g(t, \chi) - \kappa(t, \varepsilon^\alpha) - \frac{1}{2} \gamma \sigma_c^2 - \frac{\delta}{1 - 1/\psi} \right] F \\
& \left. + \pi_\tau(\tau) \left[F(t, y, \tau + \theta_\tau(\tau)) - F \right] + \delta \theta \chi^{1-1/\psi} D(\tau)^{1-1/\psi} F^{1-1/\theta} \right\},
\end{aligned}$$

The optimal abatement strategy is given by

$$\alpha_t^* = \kappa_\alpha(t, y, \cdot)^{-1} \left(\frac{y F_y + \frac{y \eta}{y + YPI} F_\tau}{(\gamma - 1) F} \right)$$

and the optimal consumption rate satisfies

$$\delta (\chi_t^*)^{-1/\psi} F^{-1/\theta} D(\tau)^{1-1/\psi} = - \frac{\partial}{\partial \chi} g(t, \chi_t^*).$$