Balancing Growth in a Warming World*

Adrien Bilal† Stanford University
Thibault Ingrand‡ Stanford University
Diego R. Känzig§ Northwestern University

July, 2026

Abstract

This paper develops a theory of balanced growth and optimal climate policy in an integrated assessment framework. Balanced growth is possible despite climate change, but at a slower rate. Two distinct policy regimes emerge because emissions permanently raise carbon stocks. For low damages, optimal growth coincides with the laissez-faire and carbon is priced at a constant rate. For high damages, optimal policy caps temperature growth to prevent escalating losses and achieve faster economic growth than in the laissez-faire by pricing carbon at an increasing rate. We quantify our model and show that empirical damage estimates span both regimes. For low damages, optimal policy achieves a 1% welfare gain with 3 to 8°C of warming per century and no change in long-run growth. For high damages, optimal policy raises annual growth by 0.8 percentage points and achieves a 26% welfare gain by capping warming at 0.4°C per century.

JEL classification: E23, H23, O44, Q54, Q56, Q58

Keywords: Climate change, balanced growth, optimal policy, social cost of carbon, integrated assessment

*We thank Philippe Aghion, Luigi Bocola, Maarten De Ridder, Chad Jones, David Hémous, Pete Klenow, Joseba Martinez, Ishan Nath, Rick van der Ploeg, Michael Peters, Noémie Pinardon-Touati, Tommaso Porzio, Lukasz Rachel, Esteban Rossi-Hansberg, Karthik Sastry, Conor Walsh, and participants at numerous seminars and conferences for helpful comments and suggestions.

†Stanford University, CEPR and NBER. E-mail: adrienbilal@stanford.edu. Web: www.sites.google.com/site/adrienbilal

‡Stanford University. E-mail: ingrand@stanford.edu

§Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com

Introduction

Economic growth transforms societies by compounding modest annual gains into vast improvements in living standards. Much of that growth has been and is still powered by fossil fuels. As greenhouse gases released by their combustion accumulate, economic losses from climate change may slow the very growth that produced those emissions. Over the long run, growth and climate change are therefore deeply intertwined.

By contrast, climate policy is largely built on the presumption that warming will stabilize on its own. Yet, the world economy is far from running out of fossil fuels or decarbonizing its energy needs, even after sizable policy intervention. Known fossil fuel reserves embed 10°C of additional warming and keep expanding. Global emissions continue to increase every year. Is balanced growth even possible in a steadily warming world? And if so, does the interplay between growth and climate change affect the design of climate policy?

This paper develops a theory of growth and optimal climate policy in an integrated assessment model when both the economy and temperature may grow for the foreseeable future. We show that this possibility fundamentally alters the positive and normative implications of climate change because greenhouse gas accumulation follows a unit root process and economic losses can escalate. Balanced growth is possible, but climate change slows it down. Two distinct policy regimes arise depending on the severity of climate damages. Low damages imply that long-run growth coincides in the laissez-faire and optimal allocations, and optimal policy prices the carbon externality at a constant rate. High damages slow growth under laissez-faire, while optimal policy prices carbon at an increasing rate to cap temperature growth and achieve faster economic growth. Our quantitative analysis shows that existing empirical damage estimates span both regimes. For low damages, optimal policy achieves a 1% welfare gain with 3 to 8°C of warming per century and no change in long-run growth. For high damages, optimal policy raises annual growth by 0.8 percentage points and achieves a 26% welfare gain by capping warming at 0.4°C per century.

Our analysis starts from the canonical integrated assessment framework (Nordhaus, 1992, 2013; Golosov et al., 2014), which we extend with a climate system with multiple carbon sinks and temperature layers (Folini et al., 2024). A representative household consumes and saves. A representative firm produces output using capital, labor, and energy. The productivity of the firm has two components: an exogenously growing component reflecting technological progress, and an endogenous component that represents climate impacts and depends on temperatures in all layers, including the atmosphere and the upper ocean. A competitive sector extracts energy at a time-varying cost that reflects technological progress or changes in resource scarcity. Energy use generates carbon emissions with a time-varying intensity that reflects improvements in renewables.

The climate system allows for an arbitrary number of carbon sinks and temperature layers as long as they satisfy mass and thermal energy conservation. Carbon emissions are distributed into atmospheric and oceanic sinks, leading to warming in atmospheric and oceanic layers. Our key departure from standard analyses is that we do not impose the existence of a backstop technology. Instead, we allow temperatures to grow for the foreseeable future.

We first characterize balanced growth absent policy intervention. Balanced growth emerges under two conditions that link economic activity and the climate system. First, log output losses are at most linear in temperature, consistent with standard econometric formulations (Dell et al., 2012; Bilal and Känzig, 2026) and canonical integrated assessment models (Nordhaus, 1992, 2013). Second, temperatures follow Arrhenius' law and depend on log carbon concentrations, consistent with the geoscientific literature (Arrhenius, 1896; Myhre et al., 1998; Dietz et al., 2021; Barrage and Nordhaus, 2024; Folini et al., 2024).

When progress in carbon-saving technologies is sufficiently fast relative to productivity growth, the carbon intensity of energy falls faster than energy demand grows. Temperatures stabilize in the long run and the economy converges asymptotically to an economy with an effective backstop technology. In that case, long-run growth is independent of climate change. We focus on the opposite region of the parameter space where emissions keep rising over time.

When progress in carbon-saving technologies is slow relative to productivity growth, energy demand outpaces carbon savings and emissions increase over time. In that case, climate change slows down economic growth. As carbon accumulates exponentially across all sinks, temperatures rise linearly over time due to Arrhenius' law. Climate damages then offset part of productivity growth because of log-linear damages. As a result, output, energy demand and emissions all grow exponentially in equilibrium at rates that are adjusted downward due to climate damages. We derive simple formulas that show how economic growth depends on the parameters governing climate damages, the climate system, and production.

Next, we characterize the efficient allocation. The key determinant of optimal policy is the social cost of carbon: the present value of future losses per unit of emissions. On a constant growth path, the social cost of carbon is increasing in climate damages. Crucially, Arrhenius' law and the permanent accumulation of carbon in the climate system imply that the social cost of carbon is increasing in the growth rate of emissions: damages depend on the stock of atmospheric carbon, which is larger relative to the current flow of emissions when emissions grow faster.

This property implies that optimal policy features a regime switch governed by the magnitude of climate damages, technological growth rates and the carbon intensity of the economy. In each regime, a simple formula continues to link output growth to these parameters, most importantly the parameter governing climate damages.

When productivity growth remains below a critical threshold, long-run economic and temperature growth coincide in the laissez-faire and in the optimal allocation because emissions either fall or grow only moderately. Equivalently, this regime emerges when climate damages are below a threshold that depends on the same parameters. Optimal policy then prices carbon at a constant rate that only affects the level of economic activity and emissions, not their growth over time. Despite moderately growing emissions and temperatures that slow down economic growth, optimal policy in this low damage regime behaves similarly to the case in which temperatures stabilize in the long run.

Optimal policy changes fundamentally once productivity growth exceeds the critical threshold and emissions grow rapidly, or, equivalently, when climate damages are sufficiently large. The social cost of carbon can then exceed the static marginal product of energy absent starker policy intervention. To restore efficiency, the planner selects slower emissions growth, which lowers the social cost of carbon. The value of emissions growth chosen by the planner is increasing in the energy share and decreasing in climate damages. By contrast, it is independent of productivity growth: regardless of how fast the economy would otherwise grow, the planner always chooses the same growth rate of emissions. This property leads us to call this optimal policy a cap.

Thus, at high damages, the optimal allocation caps emissions and temperature growth. This choice in turn achieves strictly faster economic growth than in the laissez-faire because it avoids escalating damages due to the permanent accumulation of carbon in the atmosphere and in the oceans. Pricing carbon at an increasing rate over time decentralizes this optimal allocation. The interplay between growth and climate change delivers a cap on temperature growth despite relying on entirely standard cost-benefit analysis and without appealing to a precautionary principle.

How robust are our conclusions to alternative specifications of our framework? We examine a range of variants. First, our results hold for Constant Relative Risk Aversion (CRRA) utility functions, when population grows over time, and regardless of the details of the climate system. Second, they hold when extraction costs are paid in labor units. Third, they hold when energy stocks are dynamically managed as in Hotelling (1931), and exploration is possible at a cost that rises endogenously to reflect scarcity. Fourth, they hold when we endogenize technological growth through innovation as in Grossman and Helpman (1991) and Aghion and Howitt (1992). Fifth, they hold when we endogenize progress in carbon savings through innovation in a renewable energy sector.

The magnitude of climate damages thus determines the properties of optimal policy. But does the shape of the damage function also matter? Our baseline analysis uses a log-linear damage function. Yet, many formulations have been proposed in the literature, with varying degrees of nonlinearity and empirical support. We classify damage functions and show that the degree of nonlinearity of the logarithm (rather than the level) of output losses is the relevant metric for long-run growth because technological growth is exponential.

Nonlinear damages exacerbate our baseline results. We start with log-concave damage functions for which log output losses grow less than linearly with temperature (Nordhaus, 1992, 2013; Barrage and Nordhaus, 2024). These functions can be interpreted as a reduced-form way to capture adaptation. In this case, economic growth and the climate are independent in the long run. Technological progress eventually swamps dwindling marginal damages.

Next, we consider log-convex damage functions for which log output losses grow faster than linearly with temperature (Acemoglu et al., 2012; Golosov et al., 2014) that may capture tipping points. In that case, damages escalate too fast for balanced growth and the economy grinds to a halt. Optimal policy continues to depend on two regimes, and achieves positive instead of zero growth.[1]

Log-linear damage functions thus offer a conceptual advantage: they deliver an in-

[1]Golosov et al. (2014) specify that log output losses are linear in the stock of carbon. When paired with a climate system that satisfies Arrhenius' law, their damage function is effectively super-exponential in temperature.

formative intermediate case in which climate impacts depend continuously on damage function parameters. They also offer a practical advantage: they map directly to the empirical damage specifications most commonly estimated in the literature. We therefore use our baseline log-linear damage function in the quantitative implementation.

We quantify the gains from optimal carbon policy in a parametrization of the world economy. We use standard values for the capital and energy shares in production and the capital depreciation rate. We set the rate of time preference to 3% annually. We estimate population growth at 1.5% annually. We use the climate system parametrization from Folini et al. (2024) that matches Coupled Model Intercomparison Project (CMIP) simulations from large-scale global circulation models at high and low frequencies. We estimate progress in carbon-saving technologies and fossil fuel extraction costs by matching growth in carbon emissions relative to energy use and energy use relative to output over the last four decades. Finally, we estimate exogenous productivity growth by matching output growth over the same period.

We explore a broad set of values for the magnitude of climate damages. Modern estimates range from 3% to 25% output losses per 1°C of warming depending on the empirical methodology. For illustration, we first focus on the extremes of this range, before analyzing intermediate values.

We start by considering classic estimates based on local temperature variation (Dell et al., 2012; Burke et al., 2015; Nath et al., 2025), without assuming growth effects. For comparability, we use the local temperature estimates from Bilal and Känzig (2026), whereby a 1°C increase in atmospheric temperature implies a 3.6% reduction in productivity in the long run. In this case, the economy remains in the low damage regime: laissez-faire and optimal output growth coincide at 3.1%. Optimal carbon pricing temporarily delays emissions without affecting their long-run growth rate, so that 2100 temperatures rise by an additional 1.5°C rather than 2.2°C relative to 2020 in the laissez-faire. Thereafter, temperatures rise by 3°C to 8°C per century as ocean temperatures gradually catch up to the balanced growth path. The optimal policy leads to a 1.2% consumption-equivalent welfare gain, which reflects the modest magnitude of the damage function.

Next, we consider damage estimates based on global temperature variation, whereby a 1°C increase in atmospheric temperature implies a 25% reduction in productivity in the long run (Bilal and Känzig, 2026).[2] In that case, the economy is in the high damage

[2]These magnitudes are also consistent with local temperature variation and imposing growth effects (Burke et al., 2015).

regime. While the economy keeps growing in the laissez-faire, long-run optimal growth is 0.8 percentage points higher: 3.2% instead of 2.4%. These growth benefits only materialize after the first forty years, during which they are offset by higher energy costs. Optimal carbon pricing permanently caps the growth rate of emissions, leading to slower temperature growth in the short and long run. Thus, warming by 2100 is limited to an additional 0.4°C relative to 2020. Thereafter, temperatures rise by 0.4°C per century. The optimal policy delivers a 26% consumption-equivalent welfare gain that reflects the larger magnitude of climate damages and the resulting externality, together with the qualitatively different policy response that caps temperature growth.

Does the cap bind only at the high end of the range of empirical estimates? We find that the opposite is true: the cap starts to bind just above the lower end of the range, when damages exceed 4.5% per 1°C. In fact, the 95% confidence interval around damage estimates based on local temperature variation is [-2.3%, 11.3%]. Thus, the cap binds even for a large fraction of plausible values for damage estimates based on local temperature variation. Because the cap falls sharply with damage estimates, even moderate deviations from the lower end of the range lead to sizable reductions in warming under optimal policy. For instance, damages of 10% per 1°C imply that optimal policy caps temperature growth at 1.8°C per century in the long run for a 7.5% welfare gain, well below long-run warming of 8°C per century at the lower end of the range.

Finally, we use our framework to assess whether our conclusions are quantitatively sensitive to key assumptions. One may reasonably argue that a backstop technology will eventually become available, that fossil fuels will run out, and that warming will eventually stabilize. Does our theoretical characterization still provide a reasonable approximation in this case? The answer is yes. When we introduce an anticipated backstop technology at various horizons in our framework, the allocations with and without backstop virtually coincide until the backstop arises.

Next, we ask how the analysis changes quantitatively when we depart from a log-linear damage function but recalibrate it to the same empirical targets. The results suggest that balanced growth analysis with log-linear damages delivers robust predictions for the foreseeable future. Nonlinearities matter little up to 5°C of warming. The laissez-faire remains below this threshold until 2200, and the efficient allocation remains below it even longer. However, the shape of the damage function is much more important for longer horizons in the laissez-faire where warming exceeds 5°C. In particular, growth stalls past

2200 under a log-convex damage function.

Related literature. This paper contributes to the literature studying the interplay between climate change and the macroeconomy (Hassler et al., 2016, 2024; Bilal and Stock, 2025). Most of this literature assumes that temperatures stabilize in the long run due to some combination of a backstop technology, no technological progress, or scarce fossil fuel reserves (e.g. Nordhaus, 1992, 2013; Acemoglu et al., 2012; Golosov et al., 2014; Ploeg and Withagen, 2014; Barrage and Nordhaus, 2024; Cruz and Rossi-Hansberg, 2023). However, even if fossil fuel reserves are ultimately finite, from a practical perspective they are effectively infinite once coal is included (Hassler et al., 2016). Known fossil fuel reserves embed nearly 10°C of additional warming, and more reserves are still being discovered (Covert et al., 2016). Important energy services and industrial processes are difficult to decarbonize, while carbon capture remains costly and far from offsetting emissions at scale (Davis et al., 2018; Bui et al., 2018). Temperatures might thus keep rising for multiple centuries—longer than the world economy has known sustained growth for.

We characterize the conditions that lead to balanced growth in such a scenario.[3] Our analysis reveals that long-run balanced growth arises in the laissez-faire only when certain physical and economic scaling laws hold. We build on Myhre et al. (1998), Dietz et al. (2021), and Folini et al. (2024) to incorporate a rich climate system that matches CMIP simulations at high and low frequencies. It links atmospheric forcing—and thus, temperature—to the log carbon stock due to Arrhenius' law (Arrhenius, 1896), which we show is key for balanced growth. The damage function takes a standard functional form linking log output losses to the level of temperature that directly maps into existing empirical estimates (Dell et al., 2012; Burke et al., 2015; Nath et al., 2025; Bilal and Känzig, 2026), which we show is also key for balanced growth. We characterize how departures from log-linear damages either decouple growth and climate change, or stop growth altogether.

We study the implications of a continuously warming world for optimal climate policy and the social cost of carbon. Existing work derives optimal carbon taxes in dynamic general equilibrium climate models (Golosov et al., 2014; Barrage, 2020) and integrated assessment frameworks (Nordhaus, 2013; Barrage and Nordhaus, 2024). Our analysis builds on these insights and shows that, when temperature need not stabilize, optimal

[3] Brock and Taylor (2010) characterize balanced growth in a model with emissions but do not incorporate a climate feedback through damages that is central to our analysis.

policy features two regimes that crucially depend on climate damages because the permanent accumulation of carbon in the climate system implies escalating losses in a growing world.

1 A Model of Growth with Climate Change

We introduce a general equilibrium framework that links economic growth and climate change through carbon emissions, warming, and climate damages.

1.1 Economic Module

The economic module is a simple extension of the neoclassical growth model with energy use and a climate module, similarly to Golosov et al. (2014). Time is continuous and runs forever.

Preferences. There is a unit continuum of infinitely-lived identical households who populate the world economy. Households have CRRA flow preferences: $U(C) = \frac{C^{1-\gamma}-1}{1-\gamma}$. For ease of exposition, in the main text we focus on logarithmic utility $\gamma = 1$, and report our results under general CRRA utility in the Appendix. Labor supply is exogenous and set to $L_t = 1$. The pure rate of time preference of households is ρ. Intertemporal preferences are given by $\int_0^\infty e^{-\rho t} U(C_t) dt$.

Final good sector. Firms in the final good sector produce according to a Cobb-Douglas production function. They use capital, labor and energy, with time-dependent labor productivity Z_t:

$$Y_t = \left(K_t^\alpha (Z_t L_t)^{1-\alpha}\right)^{1-\psi} E_t^\psi \equiv F(Z_t, K_t, L_t, E_t). \tag{1}$$

The path of productivity Z_t is perfectly foreseen and endogenous to the climate module.

Energy sector. Firms in the energy sector extract energy at an exogenous time-varying unit cost χ_t paid in units of the final good. This formulation is isomorphic to one in which energy is produced with the same production function as final goods, with relative total factor productivity χ_t. We show this equivalence in Appendix A.1.

Markets. All factor and goods markets—capital, labor, energy and the final good—are perfectly competitive. The final good is the numeraire. Households earn wages w_t from the labor they supply to firms in the final good sector. Households own claims to firms and capital K_t, which they rent out to firms for production at net interest rate r_t. K_0 denotes their initial capital endowment.

Resource and budget constraints. The resource constraint of the economy is: $C_t + \dot{K}_t = Y_t - \delta K_t - \chi_t E_t$. Firms make zero profits given constant returns to scale and perfect competition, implying that profits drop out of the household budget constraint: $C_t + \dot{K}_t = w_t + r_t K_t$.

1.2 Climate Module

Emissions. Energy use E_t generates greenhouse gas emissions E_t^m. For expositional simplicity, we label greenhouse gases "carbon dioxide." Carbon intensity σ_t varies exogenously over time:

$$E_t^m = \sigma_t E_t. \tag{2}$$

In Appendix A.2, we propose a simple microfoundation for time-varying carbon intensity when energy bundles two primary energy sources: emissions-free renewables and emissions-intensive fossil fuels. In that case, technological progress in either or both sectors leads to time-varying emissions intensity. If technology progresses faster in the renewable sector, carbon intensity declines over time.

Carbon cycle. The climate module describes the dynamics of carbon concentration in the environment. We consider a state-of-the-art medium-scale climate model with multiple sinks. In our quantitative application, we will use three sinks: the atmosphere, the upper ocean and the lower ocean (Folini et al., 2024). However, our theoretical analysis accommodates any number of sinks $\mathcal{M}$. The carbon cycle ensures that mass is conserved. Carbon emissions first enter the atmosphere and then gradually mix into both ocean layers.

Let $\mathbf{M}_t$ be the vector of carbon masses across the $\mathcal{M}$ sinks at time t. The dynamics of carbon masses $\mathbf{M}_t$ are given by:

$$\dot{\mathbf{M}}_t = \mathbf{A}\mathbf{M}_t + E_t^m \mathbf{m}, \tag{3}$$

where $\mathbf{m}$ is a constant vector that encodes how emissions enter carbon sinks. For instance, in our quantitative application with three sinks, we have: $\mathbf{m} = (1, 0, 0)'$ to ensure that emissions enter through the atmosphere before mixing into the oceans.

The matrix $\mathbf{A}$ encodes the physical features of carbon mixing. Carbon conservation requires $\sum_i A_{ij} = 0$ for all j. Carbon mixing requires: $A_{ij} \geq 0$ for $i \neq j$ and that $\mathbf{A}$ be irreducible: all sinks are connected at least indirectly. Our theoretical analysis only relies on these properties. Appendix B describes the exact structure of the carbon cycle matrix $\mathbf{A}$ that we use for our quantitative analysis.

Warming. The warming module describes how carbon concentrations map into temperatures. We use a general formulation of the module in Geoffroy et al. (2013). In our quantitative application, we again follow Folini et al. (2024) and consider two temperature layers: an upper layer consisting of the atmosphere and the upper oceans, and a lower layer consisting of the deep ocean. However, our theoretical analysis accommodates any number of layers $\mathcal{T} \leq \mathcal{M}$, which need not coincide with carbon sinks.

Let $\mathbf{T}_t$ be the vector of temperatures across layers at time t. Its dynamics follow:

$$\mathbf{C}\dot{\mathbf{T}}_t = (\mathbf{D} - \mathbf{L})\mathbf{T}_t + \mathbf{F}_{eq} + \mathbf{F}_t. \tag{4}$$

The matrix $\mathbf{C}$ is diagonal, with strictly positive diagonal entries that contain the thermal capacities of each temperature layer. The thermal capacity of a layer measures how much heat it can store, and therefore how much a given temperature difference with adjacent layers warms it. For instance, the thermal capacity of the ocean is larger than that of the atmosphere.

The heat exchange matrix $\mathbf{D}$ ensures that energy is conserved when energy is exchanged between temperature layers, after accounting for their respective thermal capacities, and that all layers are at least indirectly connected. These physical properties require that $\mathbf{D}$ satisfies the same properties as the carbon exchange matrix $\mathbf{A}$: $\sum_j D_{ij} = 0$ for all i, that $D_{ij} \geq 0$ for all $i \neq j$, and that $\mathbf{D}$ be irreducible.

The heat loss matrix $\mathbf{L}$ encodes that energy is radiated back to outer space. When the entire atmosphere is a layer, then only the atmosphere radiates energy to outer space. Our formulation allows us to disaggregate the atmosphere into multiple regional layers, each potentially radiating energy back to outer space. To encode these properties, $\mathbf{L}$ is diagonal with weakly positive entries, and at least one strictly positive entry.

The vector $\mathbf{F}_{eq} + \mathbf{F}_t$ denotes total radiative forcing—the amount of energy that each layer receives. $\mathbf{F}_{eq}$ represents natural radiative forcing across all layers, for instance in the atmosphere and in the deep ocean.[4] $\mathbf{F}_t$ represents the additional forcing coming from carbon emissions in the atmosphere since the pre-industrial era. Following Myhre et al. (1998), it satisfies:

$$\mathbf{F}_t = \mathbf{H}\Big(\log \mathbf{M}_t - \log \mathbf{M}_{eq}\Big), \tag{5}$$

where $\mathbf{H}$ is a rectangular $\mathcal{T} \times \mathcal{M}$ matrix that encapsulates the greenhouse effect: it contains the semi-elasticity of radiative forcing to carbon emissions. The log function applies element-wise. $\mathbf{M}_{eq}$ denotes the carbon concentration in the atmosphere in the pre-industrial era. In what follows, we refer to the property that forcing scales with the logarithm of carbon concentrations as 'Arrhenius' law' (Arrhenius, 1896).

Without loss of generality, we order all layers with atmospheric layers first, and then ocean layers. $\mathbf{H}$ is then diagonal-rectangular and captures how much radiative forcing increases following a doubling of carbon dioxide concentrations. We denote by $\mathbf{h}$ the $\mathcal{T} \times 1$ vector that collects the diagonal entries of $\mathbf{H}$ in its left $\mathcal{T} \times \mathcal{T}$ block. For instance, $\mathbf{h}$ has positive entries for each atmospheric layer, and null entries for each deep ocean layer.

We rewrite the warming module more compactly:

$$\dot{\mathbf{T}}_t = \mathbf{B}\mathbf{T}_t + \mathbf{N}(\mathbf{F}_{eq} + \mathbf{F}_t), \tag{6}$$

with $\mathbf{B} = \mathbf{C}^{-1}(\mathbf{D} - \mathbf{L})$ and $\mathbf{N} = \mathbf{C}^{-1}$. Appendix B describes the exact structure of the heat exchange matrix $\mathbf{B}$ and the forcing matrix $\mathbf{N}$ that we use for our quantitative analysis.

1.3 Climate Damages

The economic module and the climate module are coupled through climate damages. Total factor productivity depends on an exogenous component $\overline{Z}_t$ and an endogenous component that depends on atmospheric temperature relative to its pre-industrial level:

$$Z_t = \overline{Z}_t \exp\left(-\boldsymbol{\zeta}'\left(\mathbf{T}_t - \mathbf{T}_{eq}\right)\right). \tag{7}$$

[4] Deep ocean radiative forcing partly stands in for the thermohaline circulation that cools the deep ocean through the sinking of colder, less saline water at the poles. Online Appendix F shows that the quantitative properties of our climate module are virtually unchanged once we endogenize the thermohaline circulation and recalibrate our system to match CMIP simulations.

The structural damage vector $\zeta \in \mathbb{R}^{\mathcal{T}\times 1}$ governs the sensitivity of productivity to temperature across all layers. Our formulation is consistent with empirical evidence in that it features 'level effects' (Nath et al., 2025; Bilal and Känzig, 2026). Although damages operate through level effects and do not permanently reduce the growth rate, they generate widening output losses over time in a warming economy.

In our quantitative application, we use a formulation of the damage function that incorporates persistent level effects: $Z_t = \overline{Z}_t \exp\left(-\int_0^t \zeta_s'(\mathbf{T}_{t-s} - \mathbf{T}_{eq})ds\right)$. Our theoretical results extend immediately to this case provided we formulate them using the cumulative damage function. However, this extension comes at the cost of additional notation, which we abstract from for clarity of exposition in the main text.

2 Balanced Growth in a Warming World

We now analyze whether balanced growth is possible and efficient in the presence of climate change. We start our analysis with the laissez-faire, before studying the efficient allocation.

2.1 Laissez-Faire

In the laissez-faire, i.e. the decentralized equilibrium without taxes, households maximize the present value of discounted utility from consumption taking prices $\{r_t, w_t\}_t$ as given:

$$\max_{\{C_t, K_t\}_t} \int_0^\infty e^{-\rho t} U(C_t) dt \qquad \text{subject to} \qquad C_t + \dot{K}_t = r_t K_t + w_t,$$

with K_0 given, and where we have anticipated that firms make zero profits. Firms in the final good sector maximize profits taking productivity $\{Z_t\}_t$ and prices $\{r_t, w_t, p_t\}_t$ as given:

$$\max_{K_t^D, L_t^D, E_t} F(Z_t, K_t^D, L_t^D, E_t) - (r_t + \delta)K_t^D - w_t L_t^D - p_t E_t.$$

Firms in the energy sector extract energy competitively, implying that the price of energy satisfies $p_t = \chi_t$. Capital and labor markets clear: $K_t^D = K_t$ and $L_t^D = 1$.

Emissions E_t^m are related to energy use according to (2). Carbon concentrations $\mathbf{M}_t$ evolve according to (3), temperatures $\mathbf{T}_t$ evolve according to (6) and forcing F_t follows (5).

Productivity Z_t evolves according to (7).

The equilibrium conditions in the laissez-faire are standard. The Euler equation reads: $\gamma \dot{C}_t / C_t = F_K(Z_t, K_t, L_t, E_t) - \delta - \rho$, where we denote partial derivatives with subscripts, for instance: $F_K \equiv \partial F / \partial K$. Together with perfect competition and market clearing, capital and labor demand from the final good firms implies: $r_t + \delta = F_K(Z_t, K_t, L_t, E_t)$ and $w_t = F_L(Z_t, K_t, L_t, E_t)$.

The central equilibrium condition in the laissez-faire is energy demand:

$$F_E(Z_t, K_t, L_t, E_t) = \chi_t. \tag{8}$$

Firms do not internalize that energy demand today affects productivity in the future due to carbon emissions and climate damages. They only take the current energy price into account. This is the only condition that will differ from the efficient allocation.

To analyze the emergence of balanced growth, we impose the following conditions. Technology grows at rate g: $\overline{Z}_t = Z_0 e^{gt}$. Extraction costs χ_t grow at constant rate g_χ: $\chi_t = \chi_0 e^{g_\chi t}$. Emissions intensity of energy falls at rate g_σ: $\sigma_t = \sigma_0 e^{-g_\sigma t}$. These growth rates can be positive or negative. Before stating our main results, we define a few useful concepts for reference.

Asymptotic constant growth path. In an asymptotic constant growth path, all variables except temperatures asymptotically grow at constant exponential rates. An asymptotic constant growth path is unique if the growth rates do not depend on initial conditions.

Balanced growth path. A balanced growth path is an asymptotic constant growth path in which the ratios K_t / Y_t and $\chi_t E_t / Y_t$ are constant over time.

We exclude temperatures from our definition of an asymptotic growth path because their dynamics will differ across cases. We are interested in how balanced growth depends on the parameters that govern the economy and the climate system. To that end, we define four key objects.

The first object is the energy share adjusted by the labor share, defined as:

$$\pi = \frac{\psi}{(1-\alpha)(1-\psi)}.$$

The larger π, the more important energy is in production relative to labor.

The second object is the growth rate of carbon savings absent policy intervention, defined as:

$$\omega = g_\sigma + (1+\pi)g_\chi.$$

ω is higher if there is faster growth in carbon-saving technologies g_σ. It is also higher if there is faster growth in energy extraction costs g_χ, which lowers the carbon intensity of production.

The third object is the long-run elasticity of climate damages to the common component of carbon stocks, defined as:

$$\Delta = \boldsymbol{\zeta}'\boldsymbol{\theta}, \qquad \boldsymbol{\theta} \equiv (\mathbf{L} - \mathbf{D})^{-1}\mathbf{h}. \tag{9}$$

The vector $\boldsymbol{\theta}$ denotes the climate sensitivity across temperature layers: the long-run semi-elasticity of temperature to the common component of carbon stocks.

To gain intuition, consider a unidimensional climate system $\mathcal{M} = \mathcal{T} = 1$. In that case, $\mathbf{D} \equiv 0$, and $\mathbf{L} \equiv \ell > 0$ is a scalar that captures energy losses to outer space. The climate sensitivity is then $\theta = \frac{h}{\ell}$, and the elasticity of damages to carbon stocks satisfies $\Delta = \frac{\zeta h}{\ell}$. In our calibration, we use two damage functions with cumulative impacts $\zeta^{\text{local}} = 0.036$ and $\zeta^{\text{global}} = 0.25$, which span the range of empirical estimates (see Section 3.3 for details). We report calibrated values as ordered pairs, with the first entry corresponding to ζ^{local} and the second to ζ^{global}. Together with the climate system calibration from Folini et al. (2024)—$\mathbf{N} \equiv n = 0.14$, $h = 4.98$, and $\ell = 1.06$—the implied values are $\Delta = (0.17, 1.17)$.

In the unidimensional case, the long-run elasticity of damages to the carbon stock depends on three parameters. The first object is the damage parameter, ζ. When it is larger, a given temperature change has larger economic impacts. The second object is the semi-elasticity of radiative forcing to atmospheric carbon concentrations, h. When it is larger, a given amount of atmospheric carbon leads to higher radiative forcing. The third parameter ℓ represents relative energy losses to outer space. When energy losses to outer space ℓ are larger, a given amount of forcing leads to lower atmospheric temperatures because more energy is radiated back to outer space. Equation (9) generalizes this logic to a climate module of any dimensionality.[5]

The fourth object is the vector of discounted long-run elasticities of climate damages

[5] Δ is strictly positive when $\boldsymbol{\zeta} > 0$ because $\mathbf{L} - \mathbf{D}$ is an irreducible non-singular M-matrix and therefore has a strictly positive inverse.

to all carbon stocks, defined as:

$$\overline{\boldsymbol{\Delta}} = \boldsymbol{\zeta}'(\rho\mathbf{C} + \mathbf{L} - \mathbf{D})^{-1}\mathbf{H}. \tag{10}$$

With the damage functions we consider in our calibration together with the climate system calibration from Folini et al. (2024), $\sum_i \overline{\Delta}_i = (0.11, 0.75)$.

Equipped with these definitions, we state our first main result.

Proposition 1 (Growth in the Laissez-Faire)**.** *The laissez-faire admits a unique asymptotic constant growth path, which is also a balanced growth path. It satisfies:*

(i) If carbon savings grow faster than technology $g < \omega$:

Emissions fall $g_m = g - \omega < 0$, temperatures are constant, productivity grows at rate $g_Z = g$, and output grows at rate $g_Y = g - \pi g_\chi$.

(ii) If carbon savings grow as fast as technology $g = \omega$:

Emissions are constant $g_m = 0$, temperatures increase log-linearly, productivity grows at rate $g_Z = g$, and output grows at rate $g_Y = g - \pi g_\chi$.

(iii) If carbon savings grow slower than technology $g > \omega$:

Emissions grow at rate $g_m = \frac{g-\omega}{1+\Delta} > 0$, temperatures grow linearly with slopes $\dot{\mathbf{T}} = g_m \boldsymbol{\theta} > 0$, productivity grows at rate $g_Z = g - \Delta g_m$, and output grows at rate $g_Y = \frac{g+(\Delta-\pi)g_\chi+\Delta g_\sigma}{1+\Delta}$.

In addition, output growth g_Y is decreasing in damages Δ, is positive if $g_\chi + g_\sigma > 0$, and is negative if $g_\chi + g_\sigma < 0$ and $\Delta > \left|\frac{g-\omega}{g_\chi+g_\sigma}\right| - 1$.

Proof. See Appendix A.4. □

Proposition 1 continues to hold under any CRRA utility function; the proof does not rely on logarithmic utility. The first case in Proposition 1 arises when carbon savings grow faster than energy demand. Even though the economy would demand more energy over time due to rising productivity, two forces blunt the effect of rising energy demand on emissions. The first is reductions in emissions intensity that mechanically slow the accumulation of atmospheric carbon. The second is rising extraction costs and thus energy prices, which lower energy demand enough if the energy share is high.

This first case is asymptotically equivalent to an economy with an exogenous backstop technology. Under these assumptions, the carbon intensity of production drops to zero, so that emissions fall over time. As a result, atmospheric carbon converges to a constant, implying that temperatures stabilize. Consequently, output growth is independent of climate damages. This case provides a useful benchmark as our setting nests the common backstop-technology assumption, asymptotically.

The second case in Proposition 1 is the knife-edge situation in which carbon savings grow exactly as fast as production technology. Emissions are constant, and so temperatures increase log-linearly over time. This slow increase is dominated by exponential exogenous productivity growth in the long run. From an economic perspective, this case resembles the first one, as climate damages do not affect long-run economic activity.

These conclusions fail in the last case when carbon savings grow more slowly than exogenous productivity $\omega < g$. Over time, production requires more energy. The price of energy is not rising fast enough to blunt energy demand, and the emission intensity of energy is growing too slowly. A constant energy expenditure share implies that emissions rise exponentially over time. The law of motion of the carbon stocks implies that atmospheric carbon grows exponentially as well. Temperatures grow linearly because of Arrhenius' law.

Given the log-linear damage function, linearly rising temperatures drag down productivity growth proportionally to the long-run elasticity of damages to emissions, Δ, and to the growth rate of emissions. This drag on productivity growth feeds back to output growth, which is now weaker if damages Δ are larger.

In fact, Proposition 1 reveals that long-run growth can be nonpositive: the carbon externality can cause immiserizing growth. When extraction costs fall rapidly over time, output would grow faster absent climate damages. With climate damages, declines in energy costs can lead to an adverse feedback loop. Demand for energy grows so fast that carbon accumulates rapidly in the atmosphere. When damages Δ are large relative to technological progress, the economic impact of rising temperatures dominates and degrowth emerges.

2.2 Efficient Allocation

In contrast to the laissez-faire, the social planner internalizes the global carbon externality and its effect on productivity. They solve:

$$\max_{\{C_t, K_t, E_t\}_t} \int_0^\infty e^{-\rho t} U(C_t) dt$$
$$\text{subject to } C_t + \dot{K}_t = F(Z_t, K_t, L_t, E_t) - \chi_t E_t$$
$$\text{and to equations (2), (3), (6), (5), (7).}$$

As we show in Appendix A.5, the only difference between the decentralized equilibrium and the planner's allocation arises in energy use. The planner's allocation satisfies:

$$F_E(Z_t, K_t, L_t, E_t) = \chi_t + \sigma_t \tau_t, \qquad \text{with } \tau_t = \frac{\boldsymbol{\mu}_t' \mathbf{m}}{u'(C_t)}. \tag{11}$$

The planner internalizes how energy use today affects productivity today and in the future through the climate system and climate damages.

τ_t represents the social cost of carbon: the present value of all future damages associated with emitting one ton of carbon at date t. The social cost of carbon τ_t increases with the planner's shadow value of an additional unit of carbon in each sink, $\boldsymbol{\mu}_t$, and is expressed in numeraire units using the marginal utility of consumption. The social cost of carbon enters optimal energy demand after adjustment by the carbon intensity of energy: higher carbon intensity implies larger future damages from energy use today. Efficient energy demand (11) immediately implies that τ_t is the carbon tax that decentralizes the efficient allocation—and $\sigma_t \tau_t$ is the corresponding energy tax.

Whether the dynamics of the efficient allocation differ qualitatively from the laissez-faire depends on the social cost of carbon τ_t. In principle, it could grow faster or more slowly than the extraction cost χ_t. Depending on which component dominates in the long run, the efficient allocation may depart markedly from the laissez-faire.

Our second main result highlights that these considerations lead to two distinct policy regimes. When necessary, we denote outcomes in the laissez-faire with a superscript LF and in the efficient allocation with a superscript SP.

Proposition 2 (Efficient Growth)**.** *The efficient allocation admits a unique asymptotic constant growth path. There exists a threshold $\overline{g}_m \in (0, +\infty]$, such that on the asymptotic constant growth* *path:*

(i) If carbon savings grow faster than technology net of damages $g \leq \omega + (1+\Delta)\overline{g}_m$:

Growth rates coincide with those in the laissez-faire from Proposition 1, and the allocation is a balanced growth path. On the growth path, the optimal energy tax that decentralizes the efficient allocation remains bounded relative to extraction costs, and is a constant proportion of extraction costs when emissions grow.

(ii) If carbon savings grow more slowly than technology net of damages $g > \omega + (1+\Delta)\overline{g}_m$:

Emissions grow at rate $g_m^{SP} = \overline{g}_m < g_m^{LF}$, temperatures grow linearly with slopes $\dot{\mathbf{T}}^{SP} = \overline{g}_m \boldsymbol{\theta} < \dot{\mathbf{T}}^{LF}$, productivity grows at rate $g_Z^{SP} = g - \Delta\overline{g}_m$, and output grows at rate $g_Y^{SP} = \frac{g+(\pi-\Delta)\overline{g}_m+\pi g_\sigma}{1+\pi} > g_Y^{LF}$. On the growth path, the optimal energy tax that decentralizes the efficient allocation grows relative to the extraction cost.

In addition, $\overline{g}_m$ is finite if and only if $\sum_i \overline{\Delta}_i > \pi$, is strictly decreasing in each $\frac{\overline{\Delta}_i}{\pi}$, and is independent of g.

Proof. See Appendix A.5. □

Proposition 4 in Appendix A.5 extends the above proposition to the CRRA case.

Proposition 2 reveals that the structure of efficient growth paths fundamentally differs depending on parameters. When emissions fall over time $g < \omega$, growth rates coincide with the laissez-faire. This region of the parameter space is the analogue of a backstop technology, in which temperatures stabilize on their own in the long run. In that case, optimal policy only corrects the level of emissions, but not its long-run growth rate.

In the intermediate region, emissions grow moderately over time: $\omega < g < \omega + (1+\Delta)\overline{g}_m$. Long-run growth rates continue to coincide with the laissez-faire. Despite climate change slowing down growth, the planner continues to solely correct the level of energy use. Carbon is not accumulating fast enough for temperature to rise rapidly enough to justify altering the long-run path of the economy.

An entirely different growth regime emerges in configurations in which emissions grow rapidly in the laissez-faire: $g > \omega + (1+\Delta)\overline{g}_m$. To understand when and why, consider a unidimensional climate system $\mathcal{M} = \mathcal{T} = 1$. On a constant growth path, the social cost of carbon satisfies:

$$\sigma_t \tau_t = \frac{\overline{\Delta}}{\pi} \frac{g_m}{\rho + g_m} \psi \frac{Y_t}{E_t}. \tag{12}$$

When the climate system is higher-dimensional, the social cost of carbon continues to be proportional to the marginal product of energy $\psi Y_t/E_t$, with a proportionality factor that aggregates the multidimensional climate dynamics, is increasing in g_m, and follows from equation (16) in Appendix A.5.

The social cost of carbon increases with climate damages and with the climate sensitivity embedded in $\overline{\Delta}$. Similarly to the marginal benefit from energy use, the social cost of carbon scales with current economic activity. The social cost of carbon in equation (12) is reminiscent of the formula in Golosov et al. (2014), but features three important differences. The first difference is that our formula holds without additional assumptions on the savings rate, as the savings rate is always constant on the constant growth path.

The second difference is that the social cost of carbon is a constant proportion of the marginal product of energy $\psi Y_t/E_t$, rather than a constant proportion of output. This difference arises because our damage function is specified as a function of temperature instead of the carbon stock, and because temperature rises log-linearly with carbon in the long run due to Arrhenius' law, not linearly: $T_t = \theta \log M_t + o(\log M_t)$. A linear relationship can be a valid approximation for small short-run changes in carbon concentrations. However, Arrhenius' law corresponds to physical laws for larger changes such as a doubling of concentrations in the long run.

The structure of our damage function is consistent with empirical specifications used in the literature (Dell et al., 2012; Burke et al., 2015; Bilal and Känzig, 2026). By contrast and given Arrhenius' law, the specification in Golosov et al. (2014) would be equivalent to a super-exponential damage function: $Z_t = Z_0 \exp\left(gt - \zeta \exp(T_t/\theta)\right)$. This specification implicitly imposes more convex damages than existing empirical estimates and, as we show in Section 2.4, is incompatible with long-run growth.

This second difference due to Arrhenius' law is key for the long-run behavior of optimal carbon and energy pricing. Energy demand (11) immediately implies that, in the regime in which laissez-faire growth coincides with the planner's, the optimal energy tax is a constant fraction of extraction costs. By contrast, a super-exponential specification would imply that the energy tax becomes arbitrarily large relative to extraction costs.

The third and critical difference is that the social cost of carbon is increasing in the growth rate of emissions through the ratio $\frac{g_m}{\rho+g_m}$. Faster emissions growth has two effects. First, it causes carbon—and thus, economic losses—to accumulate faster relative to current energy use. The marginal damage from a marginal unit of energy today is decreasing in the carbon stock M_t due to Arrhenius' law, but is then larger relative to total energy use E_t. We express the social cost of carbon τ_t with respect to energy use E_t rather than the carbon stock M_t because it corresponds to the relevant units for the marginal product of energy $F_E = \psi \frac{Y_t}{E_t}$. The numerator g_m encodes this force. Second, faster emissions growth lowers the marginal warming potential of future emissions, again through Arrhenius' law. Because this channel operates intertemporally on future emissions, it is additionally discounted: the denominator adds $\rho + g_m$.

On net, the social cost of carbon is increasing in the growth rate of emissions g_m. When emissions growth g_m is large enough and climate damages $\overline{\boldsymbol{\Delta}}$ are large relative to the energy share π, the social cost of energy $\sigma_t \tau_t$ can exceed the marginal product of energy F_E. This imbalance is never optimal. The planner thus selects a growth rate of emissions g_m that is low enough to equalize the social cost of energy with the marginal product of energy.

In this qualitatively different regime, the growth rate of emissions is capped by the constraint $g_m = \overline{g}_m$. In the unidimensional climate system, $\overline{g}_m$ satisfies: $1 = \frac{\overline{\Delta}}{\pi} \frac{\overline{g}_m}{\rho + \overline{g}_m}$. Of course, this equation may not always have a solution if π is large. Whether there exists a solution corresponds to the condition $\overline{\Delta} > \pi$. For the climate system, we use again the calibration from Folini et al. (2024). We set the rate of time preference net of population growth to $\rho = 0.015$, and $\psi = 0.05$. As above, calibrated values are reported as ordered pairs corresponding to the two damage functions $\zeta^{\text{local}} = 0.036$ and $\zeta^{\text{global}} = 0.25$. The resulting values satisfy $\overline{\Delta} = (0.15, 1.06)$, which exceed $\pi = 0.075$, so the cap $\overline{g}_m$ is well-defined in both cases.

To assess whether the emissions growth cap is tight or loose relative to the data, we express the cap as:

$$\overline{g}_m = \frac{\rho \pi}{\overline{\Delta} - \pi}. \tag{13}$$

The cap binds when $g_m^{\text{LF}} > \overline{g}_m$, that is:

$$\overline{\Delta} \geq \pi + \frac{\rho \pi}{g_m^{\text{LF}}}. \tag{14}$$

Given the damage and climate system values discussed above, we obtain $\overline{g}_m = (0.014, 0.001)$ in the one-dimensional example. In the data, the annual growth rate of emissions is approximately $g_m^{\text{LF}} = 0.01$. Thus, when damages are estimated from local temperature variation $\zeta^{\text{local}} = 0.036$, the emissions growth cap is unlikely to bind, and long-run growth coincides in the laissez-faire and the efficient allocation.

However, these damages are just below the threshold highlighted in equation (14). The cap binds as soon as damages exceed 0.045 per 1°C, and tightens rapidly with higher damage coefficients. In that case, and in particular when damages are estimated from global temperature variation $\zeta^{\text{global}} = 0.25$, the emissions growth rate cap binds, and long-run growth diverges between the laissez-faire and the efficient allocation.

We illustrate the differences between the laissez-faire and efficient allocations in Figure 1. Panel (a) displays the growth rate of emissions as a function of the damage parameter ζ. When damages are low, the growth rate of emissions coincides in the laissez-faire and the efficient allocation. When damages cross the threshold given by equation (14), optimal policy switches regimes. The optimal growth rate of emissions hits the cap and falls sharply after a kink.

Panel (b) displays the growth rate of output as a function of the damage parameter ζ. When damages are low, output growth coincides in the laissez-faire and the efficient allocation. When damages cross the threshold given by equation (14), output growth stabilizes in the efficient allocation, while it keeps falling with damages in the laissez-faire. The benefit from a tight cap on emissions and temperature growth is to stabilize output growth in the long run.

2.3 Extensions

How specific are these results to our particular setup? We have shown that they extend from log utility to CRRA preferences and do not depend on the details of the climate system as long as carbon and temperature dynamics are linear, and radiative forcing is log-linear in carbon concentrations. In this section we show that our results continue to hold when we consider four additional departures from our setup: population growth, extraction costs in labor units, dynamic energy resource management as in Hotelling (1931), and endogenous technological growth or progress in renewables as in Grossman and Helpman (1991) and Aghion and Howitt (1992).

Population growth. We have specified our baseline framework with a fixed population for expositional simplicity. However, it immediately accommodates population growth through a simple rescaling. Suppose that instead of staying constant, population L_t grows at rate n, so that $L_t = e^{nt}$. Including population growth requires three modifications to the

Figure 1: Laissez-Faire and Efficient Growth Regimes by Damage Coefficient ζ

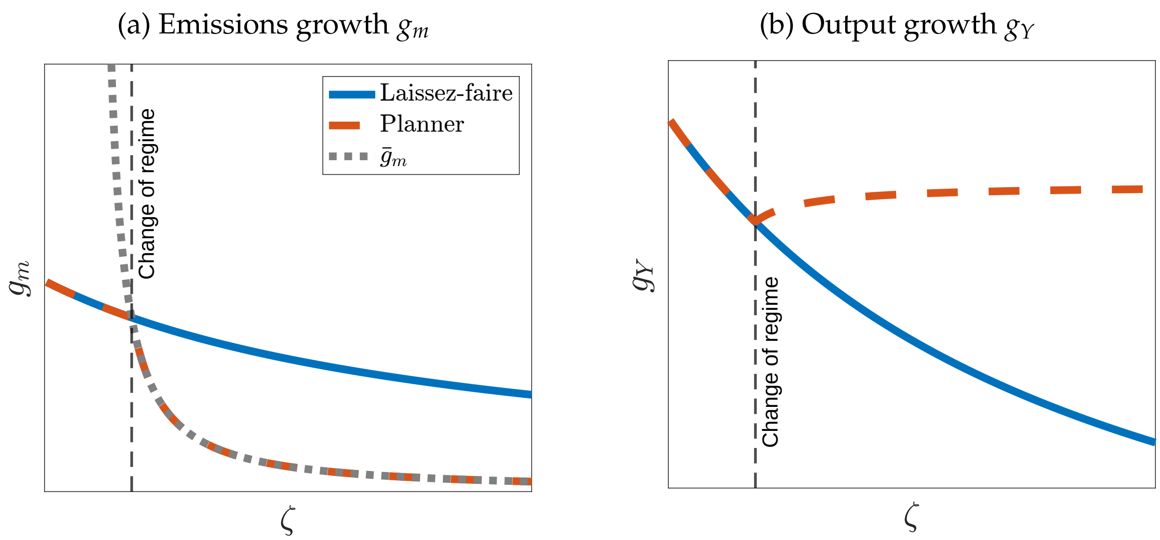

Notes: Long-run growth rates in the laissez-faire and in the efficient allocation as a function of the damage coefficient ζ, as implied by Propositions 1 and 2. Figures constructed using the calibration from Section 3. Numerical values omitted to illustrate qualitative properties. The horizontal axis varies the damage parameter ζ on the upper layer. Panel (a): growth rate of emissions g_m. Panel (b): growth rate of output g_Y. Solid blue lines: laissez-faire. Dashed red lines: efficient allocation. Dotted gray line in Panel (a): emissions growth cap $\bar{g}_m$ from equation (13). Vertical dashed lines: change of optimal policy regime, to the right of which the emissions growth cap binds.

framework. First, we shift the rate of time preference to $\rho - n$: larger future population makes dynasties effectively more patient. Second, we shift the capital depreciation rate to $\delta + n$: population growth dilutes capital per capita. Third, we shift the rate of progress in carbon savings to $g_\sigma - n$: population growth is equivalent to higher total emissions intensity of per capita emissions. With these three modifications, Propositions 1 and 2 apply. Appendix A.3 provides further details.

Extraction costs. We have assumed that producing energy requires the same bundle of inputs as production of the final good. We also explore a more demanding formulation in which energy production only requires labor, as detailed in Online Appendix E.2. Because labor is the scarce factor, this formulation implies that the cost of energy rises faster than in our benchmark specification, and thus emissions generally grow more slowly. Even then, our results continue to apply. Corollary 1 in Online Appendix E.2 shows that Propositions 1 and 2 hold in this setup.

Exhaustible resource management with exploration. We have abstracted from the dynamic management of energy stocks—for instance, fossil fuels. While the time-varying extraction cost χ_t stands in for rising scarcity, the optimal decision to extract or not affects the dynamics of energy prices, which may alter our conclusions. We consider a formulation that incorporates the dynamic management of an existing stock of energy resources as in Hotelling (1931) together with exploration at a cost that rises endogenously with scarcity in Online Appendix E.3. Corollary 2 in Online Appendix E.3 shows that Propositions 1 and 2 continue to hold in this setup.[6]

Endogenous technological growth. We have assumed that technological growth was exogenous. If technological growth is the result of private innovation decisions as in Grossman and Helpman (1991) and Aghion and Howitt (1992), optimal policy must also intervene to align them with the social optimum. This change in the long-run growth rate feeds back into emissions and may affect long-run carbon policy. We detail an extension of our framework with these features in Online Appendix E.4. We continue to obtain the same classification of growth paths within the laissez-faire, and within the efficient allocation. The only difference from our baseline results is that growth rates are generally shifted in the planner allocation because of innovation externalities—business stealing and market size—that are unrelated to the carbon externality. Corollary 3 in Online Appendix E.4 shows that Propositions 1 and 2 hold in this setup.

Endogenous progress in renewables. We have assumed that progress in renewable energy was exogenous. If progress in carbon-free energy is the result of endogenous innovation decisions, directing innovation to reduce the carbon intensity of production may complement, or even dominate, rising carbon pricing over time when emissions would otherwise escalate. We enrich our framework with endogenous growth in renewables in the tradition of Grossman and Helpman (1991) and Aghion and Howitt (1992), similarly to how we introduce it for general technological growth, in Online Appendix E.5.[7] We continue to obtain the same classification of growth paths within the laissez-faire, and within the efficient allocation. The only significant difference from our baseline results is that growth rates are generally shifted in the planner allocation because of innovation externalities that are unrelated to the carbon externality. Corollary 4 in Online Appendix E.5 shows that Propositions 1 and 2 hold in this setup.

[6]Of course, if there is a tight upper limit on how much energy can be extracted, carbon stocks and temperature eventually stabilize, and long-run growth in the laissez-faire coincides with the efficient allocation. However, if potential energy sources are unbounded and exploration is possible, temperature may grow without bound.

[7]This extension captures broadly similar forces to those in Acemoglu et al. (2012) and Hassler et al. (2021).

2.4 Nonlinear Climate Damages

An important feature of our analysis is that the damage function (7) imposes a log-linear relationship between productivity and temperature. This formulation is consistent with econometric specifications that estimate ζ using temperature variation (Dell et al., 2012; Bilal and Känzig, 2026). However, a wide range of damage function specifications have been proposed in the literature. If damages are strongly concave in temperature—perhaps because of unmodeled adaptation—they might not affect long-run growth. If damages are highly convex—perhaps because of tipping points—they might have more dramatic effects on growth.

To understand the role of the shape of the damage function, we propose a simple classification. For expositional simplicity in this section, we focus on a unidimensional climate system $\mathcal{M} = \mathcal{T} = 1$ and consider damage functions $\mathcal{D}(T)$ that are decreasing in temperature T, and such that $Z_t = \overline{Z}_t \mathcal{D}(T_t)$. For every damage function $\mathcal{D}(T)$, we denote $D(T) = -\log \mathcal{D}(T)$.

We say that a damage function $\mathcal{D}(T)$ is asymptotically *log-linear* in temperature if there exists a constant $\zeta > 0$ such that the associated log damages D satisfy: $D'(T) \to \zeta$ as T grows large. This formulation is the one we use in (7). It is the most commonly adopted specification in empirical studies of climate impacts and is also routinely used in structural models (Bilal and Känzig, 2026). Because the logarithm of output is linear in temperature, this damage function embeds nonlinearity in output levels.

We say that a damage function $\mathcal{D}(T)$ is asymptotically *log-concave* in temperature if $D'(T) \to 0$ as T grows large. This formulation implies that damages grow more slowly than log-linearly as temperature grows. It includes the specifications across a wide range of integrated assessment models (Nordhaus, 1992; Barrage and Nordhaus, 2024). In these studies, the *level* of the damage function is specified as nonlinear in temperature, for instance $\mathcal{D}(T) = \frac{1}{1+\zeta T^d}$ for $d > 0$. While this specification allows for *nonlinearities in levels*, it is sublinear in logs, and, in fact, *log-concave*: $D(T) \sim -d \log T$. Therefore, the log-concave damage functions in Nordhaus (1992) and Barrage and Nordhaus (2024) impose less nonlinearity than our baseline (7).

We say that a damage function $\mathcal{D}(T)$ is asymptotically *log-convex* in temperature if $D'(T) \to +\infty$ as T grows large. This formulation implies that damages grow much faster than log-linearly as temperature grows. It includes the specification in Golosov et al. (2014), who specify log damages as a function of the level of the carbon stock M_t.

Given Arrhenius' law, this specification is equivalent to a log damage function $D(T) = \zeta e^{T/\theta}$. Thus, the level of such a damage function is effectively super-exponential: $\mathcal{D}(T) = e^{-\zeta e^{T/\theta}}$. The log-convex damage function in Golosov et al. (2014) imposes more nonlinearity than our baseline (7).

Equipped with this classification of damage functions, we characterize how our results change when the damage function has different shapes. It is useful to define the combination of parameters that corresponds to a determinant of temperature growth in one of the cases: $\vartheta = \frac{\pi\rho(\rho+\ell n)}{n\ell}$, where $\ell \equiv \mathbf{L}$ and $n \equiv \mathbf{N}$ in the one-dimensional climate model.

Proposition 3 (Growth with a Nonlinear Damage Function)**.**

(i) *If the damage function is asymptotically log-linear, Propositions 1 and 2 apply.*

(ii) *If the damage function is asymptotically log-concave, long-run growth rates are independent of the climate and coincide in the laissez-faire (LF) and efficient (SP) allocations, and satisfy:* $g_m^i = g - \omega$, $g_Y^i = g - \pi g_\chi$, *for* $i = LF, SP$.

(iii) *If the damage function is asymptotically log-convex and if* $g < \omega$, *Propositions 1 and 2 apply.*

(iv) *If the damage function is asymptotically log-convex and if* $g > \omega$, *to leading order:*

- *In the laissez-faire:* $g_{mt} \to 0$, $g_{Yt} \to g_\chi + g_\sigma$ *and* $T_t = D^{-1}\big((g-\omega)t\big)$ *as* $t \to +\infty$.
- *In the efficient allocation:*
  - *If* $g < \omega + \vartheta$, *long-run growth rates of emissions, output and temperature coincide with the laissez-faire.*
  - *If* $g > \omega + \vartheta$, $g_{mt} \to 0$, $g_{Yt} \to g_\chi + g_\sigma + \frac{g-\omega-\vartheta}{1+\pi}$ *and* $T_t = D^{-1}(\vartheta t)$ *as* $t \to +\infty$.

Proof. See Appendix A.6. □

Proposition 3 reveals that log-linear damage functions as in our baseline (7) are the only class of damage functions for which long-run growth rates depend continuously on damages, among log-linear, log-concave and log-convex damage functions.

For log-concave damage functions as in Nordhaus (1992) and Barrage and Nordhaus (2024), climate impacts only affect the level of economic activity. The growth rate of economic activity is independent of the climate, because technological growth outpaces climate damages that grow more slowly. As a result, the laissez-faire and the efficient allocation share the same long-run growth rates. This configuration can be interpreted as one in which (unmodeled) adaptation may mute an ever-growing share of direct log-linear damages. The same conclusions arise for log-convex damage functions when carbon savings growth ω is larger than technological progress g so that temperatures stabilize.

For log-convex damage functions as in Golosov et al. (2014) and when $g > \omega$, climate impacts grow so large that the economy grinds to a halt. In the laissez-faire, output growth cannot outpace growth in carbon savings $g_\chi + g_\sigma$. If the latter is negative because energy extraction costs are falling rapidly, immiserizing growth emerges. Carbon emissions grow at an ever smaller rate that converges to zero exactly such that growing climate damages offset technological progress. As a result, the time profile of temperature is fully determined by the damage function: $T_t = D^{-1}((g - \omega)t)$. These formulas are similar to sending marginal damages Δ to infinity in Proposition 1, consistent with log-convex damages and growing temperature.

In the efficient allocation, we obtain a similar disjunction of cases to Proposition 2. When technological growth g is below carbon savings ω plus a threshold ϑ that reflects energy intensity and the climate system, growth rates coincide with the laissez-faire.

The second case arises when technological growth g exceeds carbon savings ω plus the critical determinant of temperature growth, ϑ. The planner then places a cap on emissions growth similarly to Proposition 2. However, instead of enforcing a constant exponential growth rate $\overline{g}_m > 0$ which would lead to exploding damages, the planner controls the rate at which emissions growth g_{mt} converges to zero to equalize marginal benefits and marginal costs of energy use. This convergence takes place at a different rate than in the laissez-faire.

By engineering this soft landing, the planner ensures that damage growth is capped at ϑ despite rising temperatures and log-convex damages. Thus, damages do not offset all of technological progress and output growth benefits from technological progress $g - \omega$. In our calibration, $\vartheta = 0.001$, implying that the planner achieves a strictly larger output growth rate than the laissez-faire in most parameter configurations. Temperature growth is once more determined by the damage function: $T_t = D^{-1}(\vartheta t)$.

Taking stock, Proposition 3 indicates that the degree of nonlinearity of the *logarithm* of the damage function is the relevant metric for long-run climate impacts, not the degree of nonlinearity of the *level* of the damage function. This result follows from the observation that baseline technological growth is log-linear—that is, exponential in levels—and thus dominates climate impacts from any log-concave damage function in the long run. Thus, even highly nonlinear damage functions in levels can imply no long-run climate impacts, while moderately nonlinear damage functions in logs can imply dramatic long-run climate impacts.

While the degree of nonlinearity of the damage function is of course an empirical question, log-linear damage functions have two advantages. The first advantage is conceptual: they deliver an informative intermediate case in which climate impacts depend continuously on damage function parameters. The second advantage is practical: they correspond to the majority of formulations that have been estimated in practice, and the weight of empirical evidence suggests a robustly non-zero log-linear component of the damage function. Therefore, we use our baseline log-linear damage function (7) in our quantitative implementation.

3 Quantification

To assess whether the regime change from Propositions 1 and 2 is quantitatively relevant, we specify the climate system and the parameters of our framework. First, we calibrate a set of parameters to standard values from the literature. Second, we parametrize the climate system. Third, we parametrize the damage function. Fourth, we estimate internally the key parameters that govern optimal policy—g_σ, g_χ and g—by matching out-of-balanced-growth transitional dynamics between 1982 and 2019. We start in 1982 to capture recent decades and exclude the oil shocks. We stop in 2019 to exclude the pandemic.

3.1 Data and External Parameters

Data. We obtain world Gross Domestic Product (GDP) and population data from the World Bank. Global mean temperature anomaly data is from the National Oceanic and Atmospheric Administration (NOAA). Carbon dioxide emissions from energy use are from the Potsdam Real-Time Integrated Model for the probabilistic Assessment of emission Paths (PRIMAP). We obtain world primary energy consumption from the Energy Institute through Our World in Data.

Preferences and population. We set the rate of time preference to $\rho = 0.03$ and use log utility $\gamma = 1$. We use the isomorphic formulation with population growth from Appendix A.3, with population growth set to its world average between 1982 and 2019: $n = 0.015$.

Production. Following standard practice, we set the capital share to $\alpha = 0.3$, the capital depreciation rate to $\delta = 0.08$ and the energy share to its world average of $\psi = 0.05$.

3.2 Climate System

We use the quantified climate system from Folini et al. (2024). It includes three carbon sinks: the atmosphere, the upper ocean, and the deep ocean. It features two temperature layers: a combined atmosphere-upper ocean layer and a deep ocean layer. We verify that the parametrization we use matches CMIP5 temperature impulse response functions from large-scale climate models. We describe our climate system in more detail in Appendix B.

We use the parametrized climate system to construct initial conditions for the rest of the estimation in 1982. We take 1850 to be the pre-industrial steady-state of the climate system. Given the parametrized matrices $\mathbf{A}, \mathbf{B}, \mathbf{N}, \mathbf{H}, \mathbf{F}_{eq}$, we calculate the steady-state carbon stocks $\mathbf{M}_{eq}$ in all sinks and temperatures $\mathbf{T}_{eq}$ in all layers. Next, we follow Folini et al. (2024) and input the observed path of emissions between 1850 and 1982. Iterating over the climate system delivers carbon stocks $\mathbf{M}_{1982}$ and temperatures $\mathbf{T}_{1982}$ that we use as initial conditions for the rest of the estimation. Figure 2 shows that this procedure delivers a path of global mean temperatures that aligns with the data—up to natural climate variability that we do not directly model in this paper.

3.3 Damage Function

We use an asymptotically log-linear damage function with persistent level effects that loads on temperature in the combined atmosphere-upper ocean layer with temperature T^A: $Z_t = \overline{Z}_t \exp\left(-\int_0^t \zeta_s (T^A_{t-s} - T^A_{eq}) ds\right)$.

We explore a broad set of values for the magnitude of cumulative climate damages $\overline{\zeta} = \int_0^\infty \zeta_s ds$. Modern estimates range from 3% to 25% output losses per 1°C of warming depending on the empirical methodology. We select the endpoints of this range by using

Figure 2: Climate Dynamics Between 1850 and 1982

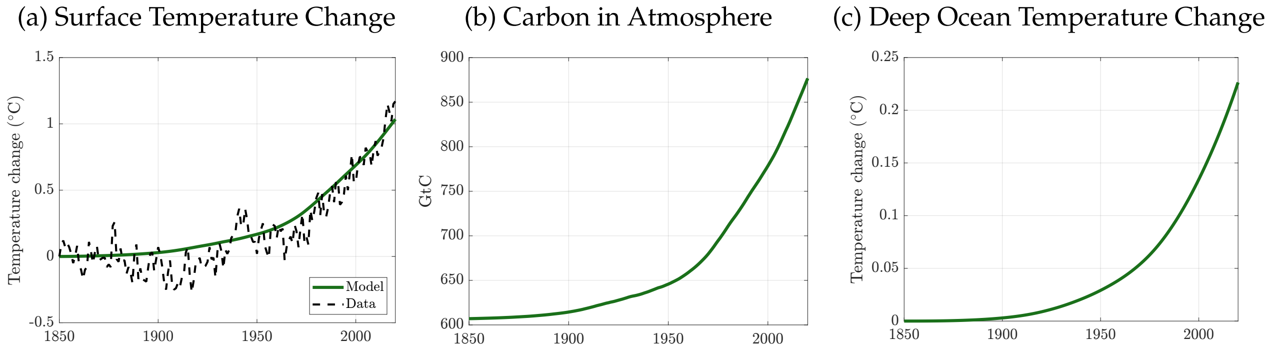

Notes: Dynamics of the climate system between 1850 and 1982, obtained by feeding the observed path of carbon emissions into the parametrized climate system of Appendix B, starting from its pre-industrial steady state in 1850. Panel (a): temperature change in the combined atmosphere and upper ocean layer. Panel (b): carbon stock in the atmosphere. Panel (c): temperature change in the deep ocean layer. Solid green lines: climate model. Dashed black lines: data. Temperatures normalized to zero in 1850.

two damage functions from Bilal and Känzig (2026) in our exercises, as their empirical strategy applies directly to our framework.

For the lower end of the range we consider the damage function ζ_s^{local}. It is estimated using annual variation in local, country-level temperature as in Dell et al. (2012) and Burke et al. (2015). This approach leads to a damage function with cumulative impact $\overline{\zeta}^{\text{local}} = \int_0^\infty \zeta_s^{\text{local}} ds = 0.036$.

For the upper end of the range we consider the damage function ζ_s^{global}. It is estimated using annual variation in global temperature directly—itself driven by phenomena such as solar cycles, El Niño or volcanic eruptions—and is a closer counterpart of the damage function in the framework. This approach leads to a damage function with cumulative impact $\overline{\zeta}^{\text{global}} = \int_0^\infty \zeta_s^{\text{global}} ds = 0.25$. It is seven times larger because global temperature fluctuations capture damaging extreme weather events such as extreme rain, wind and droughts that are not strongly correlated with idiosyncratic local temperature realizations, as well as cross-regional spillover effects.

3.4 Carbon Savings and Technological Growth Rate

We estimate the growth rate of carbon savings g_χ, g_σ and technology g internally.

Emissions intensity and extraction cost growth rates. In our framework, the growth rate of emissions to energy use, E_t^m / E_t, is equal to the growth rate of emissions intensity σ_t. Similarly, the growth rate of the extraction cost χ_t is equal to the growth rate of output per unit of energy Y_t/E_t. These relationships only depend on static equations of the model and hold away from a constant growth path. Thus, we target these average growth rates directly in the data between 1982 and 2019. Despite progress in less carbon-intensive energy generation, we estimate a relatively low value $g_\sigma = 0.003$, which, combined with population growth, is necessary to rationalize rising emissions over time. Similarly, we estimate $g_\chi = 0.012$. These estimates let us construct $\omega = (1+\pi)g_\chi + g_\sigma$.

Technological growth rate. Estimating technological growth g is more challenging. If the economy were on a constant growth path in the data, Propositions 1 and 2 would let us read off g from output growth g_Y. However, the climate system is still far from a constant growth path. Oceans take centuries to fully catch up to atmospheric temperatures, while emissions have been growing at current rates for only a few decades. Thus, the rate of warming is still below what it would be on a constant growth path in which oceanic temperatures have caught up. As a result, economic damages are growing at an increasing rate over time, and the entire economy is away from constant growth.

We address this challenge with a simulated method of moments. For each technological growth rate g, we solve the nonlinear transition of the model from its initial conditions in 1982—including the initial conditions of the climate system that we obtained in Section 3.2—to its long-run constant growth path. We search for the rate of technological growth g that allows us to match the time series of output between 1982 and 2019 to the first 38 years of the transition in the model.[8]

Figure 3 displays the results of the estimation. Panel (a) shows that output growth from 1982 to 2019 identifies technological growth g. As expected, increasing g raises the path of output upward. We infer that technological growth is $g = 0.02$ per year. Panels (b) and (c) reveal that the estimated model fits the data well on additional dimensions between 1982 and 2019. Although not surprising because we match the average growth rates of the ratio of emissions to energy, the ratio of energy to output, and the path of output, panel (b) indicates that the path of emissions from energy use in the model is close to the data at all times. Panel (c) shows that the path of global mean temperature is also close in the model and in the data between 1982 and 2019.

[8]We estimate g under the damage function ζ^{global}.

Figure 3: Estimation of g and Output and Climate Dynamics Between 1982 and 2019

(a) Output Relative to 1982 (b) Emissions from Energy (c) Surface Temperature Change

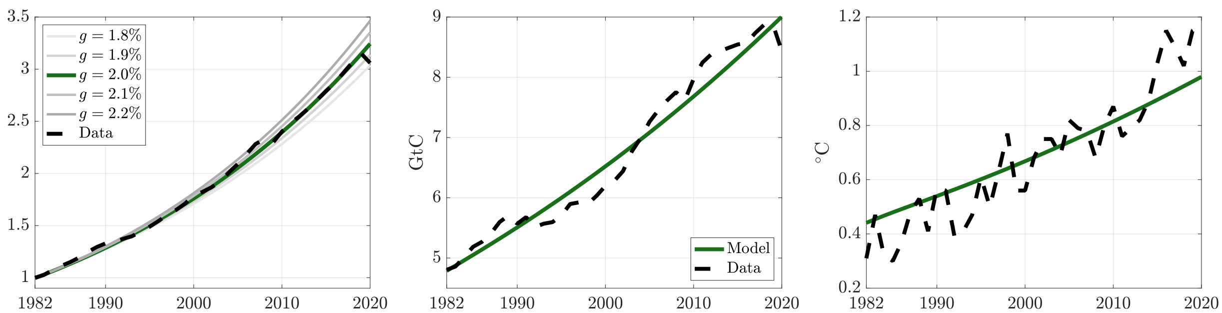

Notes: Estimation of technological growth g by simulated method of moments under the global temperature damage function ζ^{global}. The model starts from 1982 initial conditions, including the climate state from Figure 2, and is solved nonlinearly to its long-run growth path. Panel (a): world output relative to 1982 for a range of values of g, against the data. The estimated value is $g = 0.02$. Panels (b) and (c): emissions from energy use and surface temperature change at the estimated g. Solid green lines: model. Dashed black lines: data.

4 Growth and Climate Change

Equipped with the estimated model, we assess the effect of climate change and optimal policy on economic growth. We first illustrate the two policy regimes by considering the endpoints of the modern empirical range of damage functions. Next, we report our results for intermediate damages. Finally, we discuss how our results change when we introduce backstop technologies or nonlinear damage functions.

4.1 Growth and Climate Change with Low Damages

We begin our analysis under a conventional damage function ζ^{local} based on estimates using local temperature variation.[9] We start the economy from initial conditions in 2020 and compute transitional dynamics to its long-run growth path. We compare the laissez-faire and efficient allocations, and the carbon prices that implement the first-best.

Figure 4 displays transition paths under local damages for output per capita, consumption per capita and temperature. Appendix Figure 9 reports the paths of productivity, energy use per capita, emissions per capita, and carbon prices. Panel (a) shows the path of output per capita from 2020 to 2200. In the estimated model, climate damages based on local temperature variation place the economy in the regime in which the

[9] Since we estimated the model under ζ^{global}, we rescale exogenous productivity at date $t = 0$ so that initial conditions and output coincide in both specifications, making counterfactuals comparable.

Figure 4: Laissez-Faire and Efficient Transitions: Local Temperature Damage Function

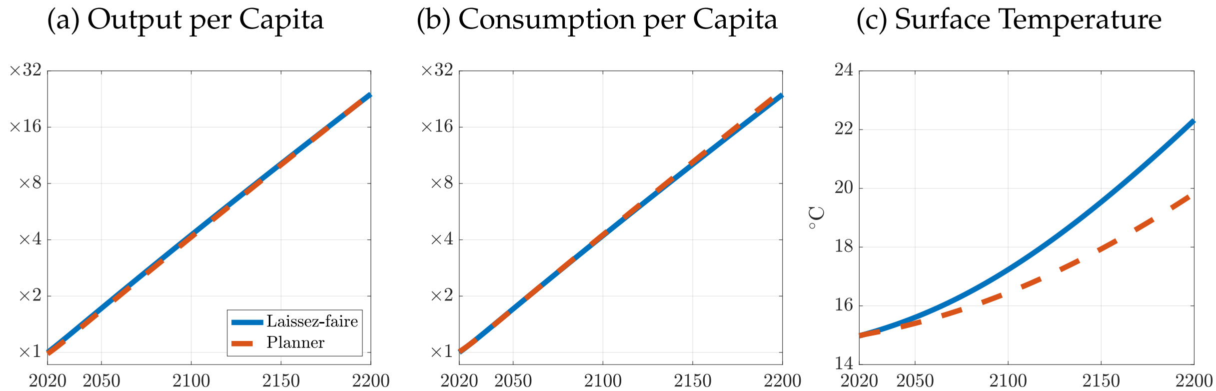

Notes: Transition paths in the laissez-faire and in the efficient allocation under the local temperature damage function ζ^{local}, estimated with local temperature variation and with cumulative impact $\int_0^\infty \zeta_s^{\text{local}} ds = 0.036$. The economy starts from 2020 initial conditions. Panel (a): log output per capita. Panel (b): log consumption per capita. Panel (c): temperature in the combined atmosphere and upper ocean layer, in °C. Solid blue lines: laissez-faire. Dashed red lines: efficient allocation. Appendix Figure 9 reports the paths of productivity, energy use per capita, emissions per capita, and carbon prices.

planner chooses the same long-run growth rate as in the laissez-faire: $g_Y^{\text{LF}} = g_Y^{\text{SP}} = 3.1\%$.

Along the transition path, efficient output falls slightly below output in the laissez-faire. This gap arises because the planner internalizes the carbon externality. The planner chooses to lower energy consumption—or decentralize it with carbon pricing—so that output initially falls below the laissez-faire. Output catches up to the level of the laissez-faire over the course of a century. Eventually, the level of efficient output exceeds the level of output in the laissez-faire because climate damages are lower.

Panel (b) indicates that despite output falling below the laissez-faire, the planner maintains consumption above the laissez-faire levels at all times due to less energy extraction and a lower need to invest in capital. While the gap in consumption levels is initially modest, it eventually opens up more visibly after multiple decades.

The consumption-equivalent welfare gain of switching from the laissez-faire to the efficient allocation is 1.2%. These gains are modest because the consumption gains are themselves moderate and largely back-loaded. This feature in turn follows from relatively minor climate damages. The damage function implies that every permanent 1°C increase in temperature leads to a 3.6% productivity loss. Given our rate of time preference, background economic growth, and the energy share, the planner can only achieve moderate gains relative to the laissez-faire.

Panel (c) reveals that even modest differences in the path of output and energy use can translate into meaningful differences in the temperature path. Under the laissez-faire, global mean temperature rises by an additional 2.2°C by 2100 relative to 2020, or 3.2°C above pre-industrial levels. By 2200, temperature is more than 7°C above 2020 levels, and thus 8°C above pre-industrial levels. By contrast, the efficient allocation reduces 2100 warming to 1.5°C above 2020 levels (2.5°C above pre-industrial levels) and 2200 warming to 4.8°C above 2020 levels (5.8°C above pre-industrial levels).

Propositions 1 and 2 indicate that temperature growth is linear on the constant growth path. While output, capital and consumption grow close to log-linearly along the transition, the temperature path is more convex. Initially, the climate system is still far from the growth path as deep ocean temperature is lagging behind atmosphere and upper ocean temperature. This initial gap between temperature in the two layers slows down temperature growth in the upper layer. Once the deep ocean temperature growth catches up, temperature growth in the atmosphere and the upper ocean accelerates and reaches its constant growth path speed. This catch-up process closes in the second half of the 22$^{\text{nd}}$ century. As this process converges, temperatures rise by 8°C per century on the balanced growth path in the laissez-faire and in the efficient allocation.

4.2 Growth and Climate Change with High Damages

We contrast the results from Section 4.1 with those under a larger damage function ζ^{global} based on estimates using global temperature variation. In that case, the cumulative damage function is nearly seven times larger and is a more comprehensive representation of the many manifestations of climate change beyond idiosyncratic, local temperature fluctuations. Propositions 1 and 2 indicate that larger climate damages can place the economy in a regime in which the efficient allocation differs substantially from the laissez-faire. Does the change of regime arise under the larger damage function based on global temperature variation?

Figure 5 displays transition paths under global temperature damages for output per capita, consumption per capita and temperature. Appendix Figure 10 reports the paths of productivity, energy use per capita, emissions per capita, and carbon prices. We find that a damage function based on global temperature variation indeed places the economy in the regime in which the efficient and laissez-faire allocations differ markedly. Panel (a) shows the path of output per capita from 2020 to 2200. The long-run growth of output is $g_Y^{\text{LF}} = 2.4\%$ in the laissez-faire, inclusive of 1.5% population growth. It falls somewhat over time in the laissez-faire, because of slow adjustments in the climate system that imply

Figure 5: Laissez-Faire and Efficient Transitions: Global Temperature Damage Function

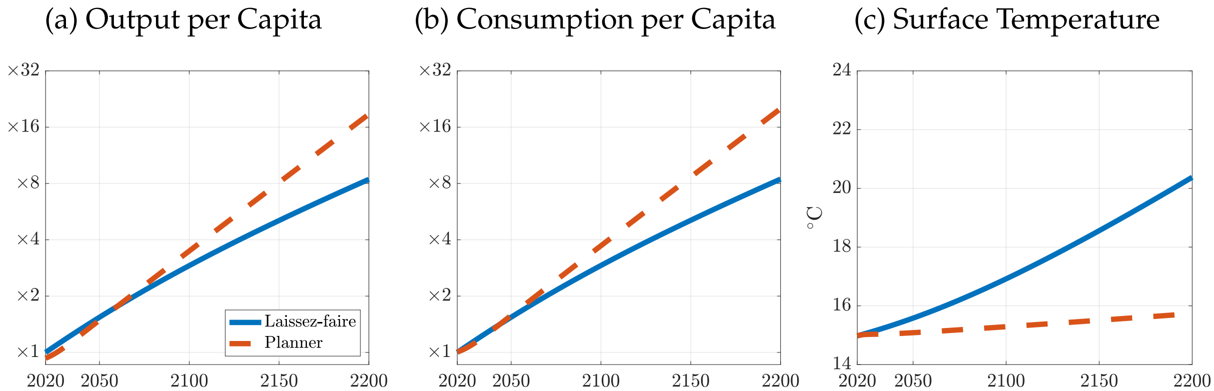

Notes: Transition paths in the laissez-faire and in the efficient allocation under the global temperature damage function ζ^{global}, estimated with global temperature variation and with cumulative impact $\int_0^\infty \zeta_s^{\text{global}} ds = 0.25$. The economy starts from 2020 initial conditions. Panel (a): log output per capita. Panel (b): log consumption per capita. Panel (c): temperature in the combined atmosphere and upper ocean layer, in °C. Solid blue lines: laissez-faire. Dashed red lines: efficient allocation. Appendix Figure 10 reports the paths of productivity, energy use per capita, emissions per capita, and carbon prices.

that damages accumulate more rapidly later in the transition. By contrast, the efficient allocation achieves a long-run output growth rate a third higher than the laissez-faire, at $g_Y^{\text{SP}} = 3.2\%$.

Along the transition path, efficient output per capita initially falls below laissez-faire output by 9%. As the planner internalizes the carbon externality, they cut back on energy use to limit the accumulation of carbon in the atmosphere. Of course, this decision is intertemporally optimal, as the planner understands that long-run reductions in climate damages are preferable to higher output today. After forty years of reduced output relative to the laissez-faire, efficient output eventually catches up because of lower climate damages. Over time, efficient output outpaces laissez-faire output markedly. In 2100, output is 20% larger in the efficient allocation than in the laissez-faire. In 2200, this ratio compounds to 123% due to the permanent gap in growth rates.

Panel (b) indicates that the initial fall in output leads to smaller short-run consumption losses of 3%: investment absorbs a large fraction of the decline in output due to less energy extraction and lower incentives to save. However, consumption catches up to the laissez-faire by 2050 and eventually substantially exceeds it. By 2200, consumption is more than 138% larger in the efficient allocation than in the laissez-faire. Consumption gains are larger than output gains because lower capital requires less investment.

The consumption-equivalent welfare gains of moving from the laissez-faire to the efficient allocation are markedly larger than under the local temperature damage function. With a global temperature damage function, the welfare gain is 26.3%, or twenty-two times the gains with a local temperature damage function. These gains are larger partly because the damage function is larger, and thus there is more to gain from implementing the optimal policy. Yet, gains are twenty-two times larger, while the cumulative damage function is seven times larger. Thus, the planner achieves larger *relative* gains when the damage function is larger. We unpack this observation in the next section.

Panel (c) reveals that the path of global mean temperature differs meaningfully between the laissez-faire and the efficient allocations. Under the laissez-faire and the global temperature damage function, the rise in global mean temperature is comparable to the rise under the local temperature damage function.[10] The efficient allocation limits 2100 warming to 0.4°C above 2020 levels (1.4°C above pre-industrial levels) and 2200 warming to 0.8°C above 2020 levels (1.8°C above pre-industrial levels).

Our analysis uncovers that the efficient allocation imposes a warming cap of 0.4°C per century, or 0.04°C per decade. This cap is twenty times tighter than the rate of warming in the laissez-faire, which reaches 8°C per century in the long run, or 0.8°C per decade. This cap arises despite our analysis relying on entirely standard cost-benefit tools and without appealing to a precautionary principle. The cap emerges as an optimal policy because carbon is a unit root process whose accumulation leads the dynamic marginal cost of energy use to exceed its static marginal product when emissions grow fast.

4.3 Climate Policy and Welfare Gains by Climate Damages

Does the optimal policy regime switch occur only at the high end of the empirical range of damage functions? To answer this question, we consider climate damages that span the entire range of modern empirical estimates. We calculate optimal policy and welfare gains for damage functions that range from no damages at all to global temperature damages, including local temperature damages, by scaling the damage function to change its cumulative impact.[11] We start by asking whether a binding temperature cap is specific to

[10]It is slightly lower because larger damages reduce output and therefore emissions. Yet, this difference in the level of flow emissions is quantitatively small compared to the accumulating stock of carbon driven by growing output and emissions.

[11]The persistence of the local temperature damage function differs slightly from the persistence of the global temperature damage function. In our comparisons in this section, we refer to the 'local temperature damage function' as the one that features the same persistence as the global temperature damage function, but is scaled so that its cumulative sum coincides with that from the baseline local temperature damage

the high end of damage estimates, or whether it may apply more broadly.

Figure 6(a) displays the rate of warming on the constant growth path as a function of the magnitude of climate damages. Consistent with Section 4.1, the laissez-faire and the efficient allocation allow for 8°C of warming per century in the long run when damages are low at 3.6% per 1°C. However, the figure reveals that the cap starts to bind for moderately larger damages that exceed 4.5% per 1°C. In fact, the 95% confidence interval around damage estimates based on local temperature variation is [-2.3%, 11.3%]. Thus, the cap binds even for a large fraction of plausible values for damage estimates based on local temperature variation.

The temperature growth cap falls steeply as damages increase, consistent with Figure 1. For instance, for intermediate damages of 10% per 1°C, optimal policy caps temperature growth at 1.8°C per century in the long run, well below long-run warming of 8°C per century when damages are at the lower end of the range. These results suggest that the optimal policy regime that involves a cap on temperature growth is not specific to the high end of damage estimates, but rather applies to a broad range of plausible damage functions.

These distinct policy regimes correspond to different welfare gains. Figure 6(b) displays welfare gains from switching from the laissez-faire to the efficient allocation as a function of the magnitude of climate damages. We compare these welfare gains to welfare gains in a scenario in which temperature exogenously stops increasing in 2020. This latter scenario is of course infeasible in our framework, but provides a useful benchmark.

The welfare gains from the optimal policy are increasing with the magnitude of climate impacts. Larger damages imply that there is more to gain from correcting the carbon externality.

Welfare gains from holding temperatures constant at 2020 values scale nearly proportionally with climate impacts. Climate damages are a constant proportion of output, and thus eliminating them increases output, consumption and welfare by a constant proportion. These gains represent the unconditional losses from the carbon externality.

By contrast, welfare gains from implementing the efficient allocation are convex in climate impacts. At low climate impacts, welfare gains remain close to zero. At large climate impacts, welfare gains rise nearly linearly. Why does the planner recover only a small proportion of unconditional losses from climate change at low damages, but a

function.

Figure 6: Long-Run Warming and Welfare Gains by Climate Damages

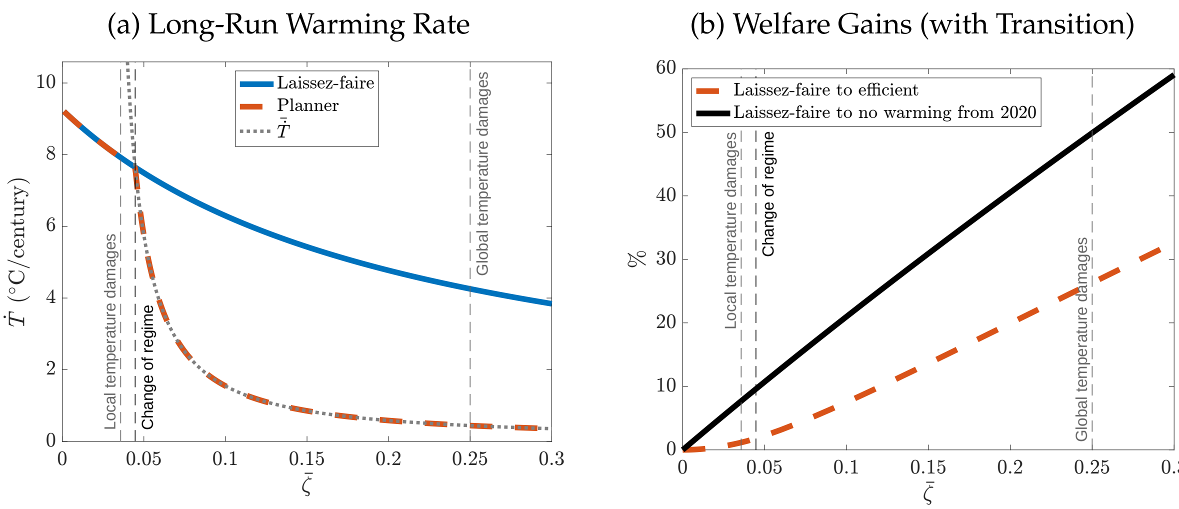

Notes: Long-run warming and welfare gains as a function of the magnitude of climate damages. The horizontal axis varies the cumulative damage function $\bar{\zeta} = \int_0^\infty \zeta_s ds$ between 0 and 0.3 by scaling the global damage function while holding its persistence profile constant. Panel (a): rate of warming on the constant growth path, in °C per century, in the laissez-faire and in the efficient allocation. Solid blue line: laissez-faire. Dashed red line: efficient allocation. Dotted gray line: warming cap implied by the emissions growth cap $\bar{g}_m$ from equation (13). Panel (b): welfare gains, measured as consumption-equivalent variations relative to the laissez-faire, from the efficient allocation and from a counterfactual in which temperature exogenously stops increasing in 2020. Solid black line: welfare gains if temperature exogenously stabilized in the initial period. Dashed red line: welfare gains from implementing the efficient allocation. Vertical lines: local temperature damage calibration (0.036), change of optimal policy regime (0.045), and global temperature damage calibration (0.25).

larger proportion at high damages? To understand these patterns, recall that the efficient allocation maximizes welfare at all values of damages, including no damages at all. When damages are low, the economy is close to the no-damage efficient allocation. Any reallocation of consumption and energy use leads to second-order gains, while direct losses due to damages are first-order: the planner has only a limited ability to offset climate change impacts. By contrast, when damages are large, the economy is far away from the no-damage efficient allocation. Second-order gains from reallocating consumption and energy use can become comparable to direct losses, and the planner manages to recover close to a constant proportion of climate change impacts.

Having established that the temperature cap binds across a wide range of damage estimates and can lead to sizable welfare gains, we examine the robustness of our results in the quantitative framework. The final two sections assess whether the arrival of a backstop technology or departures from log-linear damage functions affect our conclusions.

4.4 Backstop Technology

So far our analysis abstracts from technological breakthroughs that would stabilize temperatures at near-zero cost. While the likelihood of such breakthroughs ever occurring is necessarily uncertain, one may reasonably argue that a backstop technology will become available, that fossil fuels will run out, and that warming will eventually stabilize one way or another. Does our theoretical characterization in which temperatures keep rising provide a reasonable approximation to the allocation if a backstop were available? We answer this question by introducing a backstop technology at various horizons in our quantified model. The future arrival of the backstop technology is anticipated from 2020 onward.

When the backstop emerges, temperatures exogenously stabilize, and remain constant thereafter. We consider the high damage case, in which the backstop is likely to make the largest difference. We explore various dates at which the backstop arises: 2050, 2100, 2200 and 2500. Figure 7 displays output per capita in the laissez-faire and the efficient allocation, and Appendix Figure 11 reports other outcomes.

The allocations with and without a backstop technology are virtually identical prior to the arrival of the technology. This conclusion holds for all backstop arrival dates and for the laissez-faire and the efficient allocation. There are some anticipation effects in the efficient allocation, where the shadow cost of carbon falls as the backstop arrival nears. Quantitatively, however, these anticipation effects are small.

Of course, once the backstop technology has arrived, the allocations differ meaningfully from those without the backstop. Once temperatures stabilize, climate damages stop growing. The economy then switches to a different growth rate, and the allocations start drifting apart. The divergence is less pronounced in the efficient allocation, where the planner achieves a growth rate that is closer to the one without damages.

Why are there almost no anticipation effects? In the laissez-faire, energy demand (8) is static: firms fail to internalize the dynamic implications of carbon accumulation, implying that the future arrival of the backstop does not change decisions much beyond its implications for future growth. In the efficient allocation, the planner internalizes that carbon follows a unit root process. Any additional ton of emissions remains in the carbon cycle, depressing output forever. Thus, it remains optimal to cap temperature growth regardless of the arrival of the backstop.

Figure 7: Backstop Technologies: Output per Capita

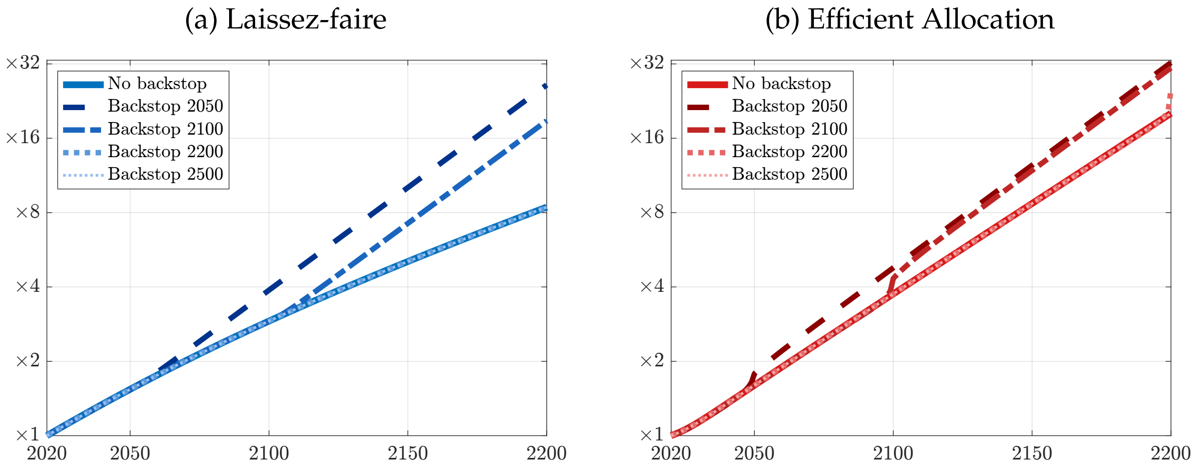

Notes: Log change since 2020 in output per capita when a backstop technology emerges at various horizons. Damage function calibrated to global temperature variation ζ^{global}. Temperature exogenously stabilizes at the level reached at the date at which the backstop arrives. Panel (a): laissez-faire. Panel (b): efficient allocation. Solid lines: baseline allocation without backstop. Dashed, dash-dotted, dotted and thin dotted lines: backstop arriving in 2050, 2100, 2200 and 2500.

4.5 Nonlinear Damages

So far our quantitative analysis has relied on a log-linear damage function. Proposition 3 shows that nonlinear damages exacerbate our baseline results: long-run growth either remains independent of climate change for log-concave damages, or grinds to a halt for log-convex damages. Empirically, it is challenging to distinguish between these three cases: plausibly exogenous historical fluctuations in temperature are too small to conclusively rule in or rule out specific functional forms.

We explore how our results change when we use alternative damage functions. First, we consider a quadratic damage function in levels as in Nordhaus (1992) and Barrage and Nordhaus (2024). This damage function, while convex in levels, is log-concave. Second, we consider a super-exponential damage function, which is log-convex and corresponds to the one used in Golosov et al. (2014) after incorporating Arrhenius' law.

We recalibrate each damage function so that its marginal damages coincide with those of the log-linear damage function at 1°C of warming above pre-industrial levels, the historical warming at which the empirical estimates are identified. We then solve the model quantitatively, and compare the resulting dynamics between damage functions. Figure 8 displays output per capita in the laissez-faire and in the efficient allocation. Appendix Figure 13 depicts additional outcomes. For this exercise we report results up to year 3000

Figure 8: Nonlinear Damage Functions: Output per Capita

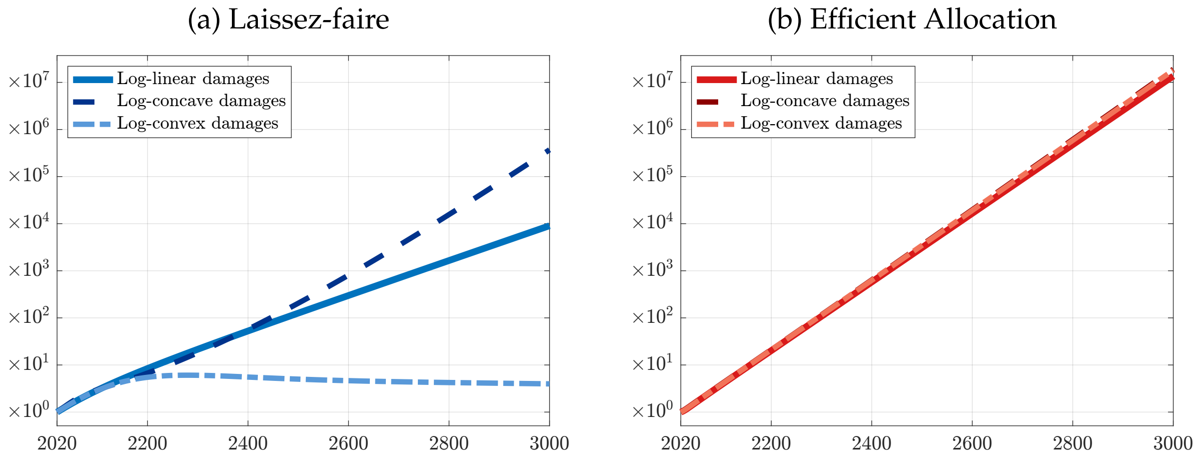

Notes: Log change since 2020 in output per capita under alternative damage functions, global calibration, between 2020 and 3000. The log-concave (Nordhaus, 1992; Barrage and Nordhaus, 2024) and log-convex (Golosov et al., 2014) damage functions are recalibrated so that marginal damages coincide with the log-linear damage function at 1°C of warming above pre-industrial levels, the historical warming at which the empirical estimates are identified. Panel (a): laissez-faire. Panel (b): efficient allocation. Solid lines: log-linear damages. Dashed lines: log-concave damages. Dash-dotted lines: log-convex damages.

to illustrate the role of nonlinearities in the very long run.

Three conclusions emerge. First, for the foreseeable future—until 2200—the particular functional form of the damage function does not strongly affect the results. Provided the damage function is calibrated to the same empirical target, warming up to 5°C does not trigger strong nonlinearities: Appendix Figure 12 shows that marginal damages remain of the same order of magnitude in that range. Of course, the log-linear damage function delivers a transparent growth path, along which the characterization of Section 2 is possible. We view this property as a theoretical advantage.

Second, in the very long run, the shape of the damage function matters substantially more. After 5°C of warming, marginal damages in the log-convex damage function rise steeply, stifling economic growth after 2200. By contrast, marginal damages in the log-concave damage function fall. In that case, in the very long run, economic growth remains independent of climate change, and growth occurs faster than with log-linear damages.

Third, the efficient allocation and the underlying optimal policy remain largely unchanged over any reasonable time horizon. The planner internalizes climate damages, regardless of their particular functional form. Thus, they tightly limit warming by imposing an exact or approximate temperature growth cap. Long-run dynamics remain similar because temperature rises little, limiting the role of nonlinearities.

5 Conclusion

This paper studies long-run growth and optimal climate policy in an integrated assessment framework in which temperatures may continue to increase for the foreseeable future. Because carbon accumulates permanently, climate damages affect long-run growth. Balanced growth remains possible in the laissez-faire, but climate damages slow it down. Optimal policy is characterized by a regime switch. When damages are below a threshold, the planner corrects levels but not long-run growth rates, and carbon is priced at a constant rate relative to extraction costs. When damages are above that threshold, the planner places an endogenous cap on emissions growth, which slows temperature growth and raises long-run economic growth relative to the laissez-faire.

Quantitatively, this regime switch is key for welfare and for the long-run trajectory of the economy. Under conventional damage estimates based on local temperature variation, optimal policy delivers modest welfare gains and largely preserves laissez-faire growth, even though it lowers the temperature path. However, even for moderate deviations from these estimates, and for larger damage estimates based on global temperature variation, optimal policy induces a qualitatively different transition: it front-loads mitigation, compresses long-run warming, and increases long-run growth by limiting escalating productivity losses from cumulative carbon. As a result, welfare gains rise more than proportionally with damages. Taking the very long run seriously shows that standard cost-benefit analysis can imply tight temperature-growth targets, without invoking the precautionary principle.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012). "The Environment and Directed Technical Change". *American Economic Review* 102.1, 131–66.

Aghion, Philippe and Peter Howitt (1992). "A Model of Growth through Creative Destruction". *Econometrica* 60.2, pp. 323–351.

Arrhenius, Svante (1896). "On the Influence of Carbonic Acid in the Air upon the Temperature of the Ground". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*. 5th ser. 41.251, pp. 237–276.

Barrage, Lint (2020). "Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy". *The Review of Economic Studies* 87.1, pp. 1–39.

Barrage, Lint and William D. Nordhaus (2024). "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model". *Proceedings of the National Academy of Sciences of the United States of America* 121.13, e2312030121.

Bilal, Adrien and Diego R Känzig (2026). "The Macroeconomic Impact of Climate Change: Global Versus Local Temperature". *The Quarterly Journal of Economics* 141.2, pp. 889–944.

Bilal, Adrien and James H Stock (2025). *A Guide to Macroeconomics and Climate Change*. Working Paper 33567. National Bureau of Economic Research.

Brock, William A. and M. Scott Taylor (2010). "The Green Solow model". *Journal of Economic Growth* 15.2, pp. 127–153.

Bui, Mai, Claire S. Adjiman, André Bardow, Edward J. Anthony, Andy Boston, Solomon Brown, Paul S. Fennell, Sabine Fuss, Amparo Galindo, Laura A. Hackett, Jason P. Hallett, Howard J. Herzog, George Jackson, Jasmin Kemper, Samuel Krevor, Geoffrey C. Maitland, Michael Matuszewski, Ian S. Metcalfe, Camille Petit, Graeme Puxty, Jeffrey Reimer, David M. Reiner, Edward S. Rubin, Stuart A. Scott, Nilay Shah, Berend Smit, J. P. Martin Trusler, Paul Webley, Jennifer Wilcox, and Niall Mac Dowell (2018). "Carbon capture and storage (CCS): the way forward". *Energy & Environmental Science* 11.5, pp. 1062–1176.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (2015). "Global non-linear effect of temperature on economic production". *Nature* 527.7577, pp. 235–239.

Covert, Thomas, Michael Greenstone, and Christopher R Knittel (2016). "Will we ever stop using fossil fuels?" *Journal of Economic Perspectives* 30.1, pp. 117–138.

Cruz, José-Luis and Esteban Rossi-Hansberg (2023). "The Economic Geography of Global Warming". *The Review of Economic Studies* 91.2, pp. 899–939.

Davis, Steven J., Nathan S. Lewis, Matthew Shaner, Sonia Aggarwal, Doug Arent, Inês L. Azevedo, Sally M. Benson, Thomas Bradley, Jack Brouwer, Yet-Ming Chiang, Christopher T. M. Clack, Armond Cohen, Stephen Doig, Jae Edmonds, Paul Fennell, Christopher B. Field, Bryan Hannegan, Bri-Mathias Hodge, Martin I. Hoffert, Eric Ingersoll, Paulina Jaramillo, Klaus S. Lackner, Katharine J. Mach, Michael Mastrandrea, Joan Ogden, Per F. Peterson, Daniel L. Sanchez, Daniel Sperling, Joseph Stagner, Jessika E. Trancik, Chi-Jen Yang, and Ken Caldeira (2018). "Net-zero emissions energy systems". *Science* 360.6396, eaas9793.

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* 4.3, pp. 66–95.

Dietz, Simon, Frederick van der Ploeg, Armon Rezai, and Frank Venmans (2021). "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists* 8.5, pp. 895–921.

Folini, Doris, Aleksandra Friedl, Felix Kübler, and Simon Scheidegger (2024). "The Climate in Climate Economics". *The Review of Economic Studies*.

Geoffroy, Olivier, David Saint-Martin, Dirk J.L. Olivié, Aurore Voldoire, Gilles Bellon, and Sophie Tytéca (2013). "Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments". *Journal of climate* 26.6, pp. 1841–1857.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014). "Optimal Taxes on Fossil Fuel in General Equilibrium". *Econometrica* 82.1, pp. 41–88.

Grossman, Gene M and Elhanan Helpman (1991). "Trade, knowledge spillovers, and growth". *European Economic Review* 35.2-3, pp. 517–526.

Hassler, John, Per Krusell, and Conny Olovsson (2021). “Directed Technical Change as a Response to Natural Resource Scarcity”. *Journal of Political Economy* 129.11, pp. 3039–3072.

— (2024). “The Macroeconomics of Climate Change: Starting Points, Tentative Results, and a Way Forward”. *National Bureau of Economic Research Working Paper Series 32197*.

Hassler, John, Per Krusell, and Anthony A Smith Jr (2016). “Environmental macroeconomics”. *Handbook of macroeconomics*. Vol. 2. Elsevier, pp. 1893–2008.

Hotelling, Harold (1931). “The economics of exhaustible resources”. *Journal of Political Economy* 39.2, pp. 137–175.

Myhre, Gunnar, Eleanor J. Highwood, Keith P. Shine, and Frode Stordal (1998). “New estimates of radiative forcing due to well mixed greenhouse gases”. *Geophysical Research Letters* 25.14, pp. 2715–2718.

Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow (2025). “How Much Will Global Warming Cool Global Growth?” *National Bureau of Economic Research Working Paper Series 32761*.

Nordhaus, William D. (1992). “An optimal transition path for controlling greenhouse gases”. *Science* 258.5086, pp. 1315–1319.

— (2013). “Integrated economic and climate modeling”. *Handbook of Computable General Equilibrium Modeling*. Vol. 1. Elsevier, pp. 1069–1131.

Ploeg, Frederick van der and Cees Withagen (2014). “GROWTH, RENEWABLES, AND THE OPTIMAL CARBON TAX”. *International Economic Review* 55.1, pp. 283–311. (Visited on 07/29/2026).

Appendix

A Proofs

A.1 Energy cost

In this section we show the equivalence between an energy cost denominated in units of the final good—as in the main text—and energy production using the same production function as for the final good.

Suppose that energy is produced according to: $E_t = \chi_t^{-1} F(Z_t, K_t^E, L_t^E, E_t^E)$, where E superscripts denote inputs into the production of energy, and χ_t now indexes relative total factor productivity in the energy sector. We use Y superscripts to denote inputs into production of the final good. Similarly, we have: $Y_t = F(Z_t, K_t^Y, L_t^Y, E_t^Y)$. Finally: $E_t = E_t^E + E_t^Y$.

In the laissez-faire and in the planning allocation, the marginal products of inputs are equalized across sectors because of factor mobility. Let $\kappa_t^S = K_t^S/(Z_t L_t^S)$ and $\varepsilon_t^S = E_t^S/(Z_t L_t^S)$ denote input ratios for sector S. We obtain: $F_I(1, \kappa_t^Y, 1, \varepsilon_t^Y) = F_I(1, \kappa_t^E, 1, \varepsilon_t^E)$ for each input I. Given constant returns, these conditions imply that input ratios $\kappa_t^S \equiv \kappa_t$ and $\varepsilon_t^S \equiv \varepsilon_t$ are equalized across sectors $S = Y, E$. Therefore, we define: $F_t^* = F(Z_t, \kappa_t Z_t, 1, \varepsilon_t Z_t)$ so that $Y_t = L_t^Y F_t^*$ and $E_t = L_t^E F_t^*$.

The resource constraint becomes:

$$C_t + \dot{K}_t + \delta K_t = Y_t = L_t^Y F_t^* = (L_t - L_t^E) F_t^* = \overline{Y}_t - \chi_t E_t,$$

where $\overline{Y}_t = F(Z_t, K_t, L_t, E_t)$ since input ratios are common across sectors, so that they are also equal to aggregate input ratios: $K_t/(Z_t L_t) = \kappa_t$ and $E_t/(Z_t L_t) = \varepsilon_t$. Therefore $\overline{Y}_t$ corresponds to output as if all resources were allocated to the final good. This resource constraint coincides with the one in the main text.

The household budget constraint becomes: $C_t + \dot{K}_t = r_t K_t + w_t$. By constant returns and factor mobility, the final good and energy firm problems lead to: $Y_t = (r_t + \delta) K_t^Y + w_t L_t^Y + p_t E_t^Y$ and $E_t = \chi_t^{-1}[(r_t + \delta) K_t^E + w_t L_t^E + p_t E_t^E]$. Therefore: $r_t K_t + w_t = \overline{Y}_t - p_t E_t$ as in the main text.

A.2 Microfoundation of Time-Varying Carbon Intensity

Assume that energy is a Cobb-Douglas aggregator of two energy sources: fossil fuels B_t that are emissions-intensive and renewables G_t that are emissions-free.

$$E_t = B_t^{1-\beta}\, G_t^{\beta}.$$

Production of fossil fuels and of renewables requires χ_t^B and χ_t^G units of the final good, respectively. Competitive markets imply that production of the energy bundle costs:

$$\mathcal{C}(E_t) = E_t \left(\frac{p_t^B}{1-\beta}\right)^{1-\beta} \left(\frac{p_t^G}{\beta}\right)^{\beta} = \Theta E_t (\chi_t^B)^{1-\beta} (\chi_t^G)^{\beta} = E_t \chi_t,$$

where $\Theta = \beta^{-\beta}(1-\beta)^{-(1-\beta)}$ and we define $\chi_t = \Theta(\chi_t^B)^{1-\beta}(\chi_t^G)^{\beta}$. This expression coincides with the one we use in the main text and relates the energy bundle extraction cost to the extraction cost of fossil fuels and renewables.

Expressing fossil-fuel use in the same units as emissions, we have:

$$E_t^m = B_t = (1-\beta)\frac{p_t^E}{p_t^B} E_t = (1-\beta)\Theta \left(\frac{p_t^G}{p_t^B}\right)^{\beta} E_t = (1-\beta)\Theta \left(\frac{\chi_t^G}{\chi_t^B}\right)^{\beta} E_t = \sigma_t E_t,$$

where the second equality uses optimal demand of fossil fuels by the intermediaries that bundle fossil fuels and renewables, and where we defined: $\sigma_t = (1-\beta)\Theta \left(\frac{\chi_t^G}{\chi_t^B}\right)^{\beta}$.

Finally, with our convention of a negative sign on g_σ:

$$g_\chi = (1-\beta) g_{\chi B} + \beta g_{\chi G} \qquad g_\sigma = \beta g_{\chi B} - \beta g_{\chi G} \qquad g_\chi + g_\sigma = g_{\chi B}.$$

With a Cobb-Douglas energy bundle, technical progress in renewables leads to higher fossil fuel demand because of the unit elasticity of substitution. As a result, $g_\chi + g_\sigma = g_{\chi B}$ is entirely determined by technological progress in fossil fuels.

A.3 Population Growth

In this section we show that our model can be rescaled to accommodate population growth. Assume that population L_t grows at rate n: $L_t = e^{nt}$. We denote per capita variables as before, and aggregate variables with a bar. Dynastic preferences are then:

$$\int_0^\infty e^{-\rho t} L_t U(C_t) dt = \int_0^\infty e^{-(\rho-n)t} U(C_t) dt.$$

Thus, we need to shift the rate of time preference to $\rho - n$.

The production function implies that aggregate output $\overline{Y}_t$ satisfies:

$$\overline{Y}_t = \left((L_t Z_t)^{1-\alpha} \overline{K}_t^{\alpha} \right)^{1-\psi} \overline{E}_t^{\psi} = L_t \left(Z_t^{1-\alpha} K_t^{\alpha} \right)^{1-\psi} E_t^{\psi} = L_t Y_t = L_t F(Z_t, K_t, 1, E_t).$$

The resource constraint is:

$$\overline{C}_t + \dot{\overline{K}}_t = \overline{Y}_t - \delta \overline{K}_t.$$

Using that $\dot{\overline{K}}_t = \dot{K}_t L_t + K_t L_t n$, we obtain the per capita resource constraint:

$$C_t + \dot{K}_t = Y_t - (\delta + n) K_t.$$

Thus, we need to shift the capital depreciation rate to $\delta + n$.

The budget constraint of the representative dynasty is:

$$\overline{C}_t + \dot{\overline{K}}_t = r_t \overline{K}_t + w_t.$$

As for the resource constraint:

$$C_t + \dot{K}_t = (r_t - n) K_t + w_t.$$

Thus, we need to shift the interest rate to $r_t - n$.

Emissions are:

$$\overline{E}_t^m = \overline{\sigma}_t L_t E_t = \sigma_t E_t,$$

where the growth rate of σ_t is now $g_\sigma - n$.

The firm's optimality conditions are:

$$r_t + \delta = \alpha(1-\psi) \frac{\overline{Y}_t}{\overline{K}_t} = \alpha(1-\psi) \frac{Y_t}{K_t} \qquad w_t = (1-\alpha)(1-\psi) \frac{Y_t}{L_t} \qquad p_t = \psi \frac{Y_t}{E_t}.$$

We can rewrite the capital demand equation:

$$(r_t - n) + (\delta + n) = \alpha(1-\psi)\frac{Y_t}{K_t},$$

which is then consistent with our rescaling above.

A.4 Proof of Proposition 1

A.4.1 Economic module

Consider an asymptotic growth path. Let g_Z be the asymptotic growth rate of productivity. The resource constraint requires that the growth rates of output, consumption and capital are equal: $g_Y = g_C = g_K$. The Cobb-Douglas production function implies a constant energy share, and thus: $g_\chi + g_E = g_Y$. The structure of the production function then requires: $g_Y = \alpha(1-\psi)g_Y + (1-\alpha)(1-\psi)g_Z + \psi g_E$, which we rewrite: $(1+\pi)g_Y = g_Z + \pi g_E$. Substituting out g_E and solving for g_Y yields: $g_Y = g_Z - \pi g_\chi$. From the definition of emissions, they grow asymptotically at rate: $g_m = g_E - g_\sigma = g_Y - g_\chi - g_\sigma$, where we choose the convention: $\sigma_t = \sigma_0 e^{-g_\sigma t}$.

A.4.2 Climate module

We now turn to the asymptotic behavior of the climate module. We divide the analysis into three different cases, corresponding to three different emissions regimes.

Case 1: Growing emissions $\boldsymbol{g_m > 0}$. We start from $E_t^m = \overline{m}e^{g_m t} + \mathcal{O}(1)$ on the growth path. We organize the proof in a sequence of claims.

Claim: $\mathbf{M}_t$ grows at rate g_m asymptotically.

To prove this claim, observe that the asymptotic behavior of the carbon cycle is given by the solution to: $\dot{\mathbf{M}}_t = \mathbf{A}\mathbf{M}_t + \overline{m}e^{g_m t}\,\mathbf{m} + \mathcal{O}(1)$. The explicit solution is: $\mathbf{M}_t = e^{t\mathbf{A}}\mathbf{M}_0 + \overline{m}e^{g_m t}\int_0^t e^{s(\mathbf{A}-g_m\mathbf{I})}ds\;\mathbf{m}$, where $\mathbf{I}$ denotes the identity matrix. Since $\mathbf{A}$ has a single zero eigenvalue and all other eigenvalues have strictly negative real parts, and $g_m > 0$, the homogeneous component $e^{t\mathbf{A}}\mathbf{M}_0$ converges exponentially to a constant, and the integral $\int_0^t e^{s(\mathbf{A}-g_m\mathbf{I})}ds$ converges exponentially to a constant as well. Therefore, $\mathbf{M}_t$ grows at rate g_m asymptotically.

Claim: $\mathbf{T}_t$ grows linearly at speed $g_m\boldsymbol{\theta}$ asymptotically.

To prove this claim, observe that since forcing is log-linear in atmospheric CO_2, which (asymptotically) grows at a constant rate, forcing must (asymptotically) grow linearly with a constant slope $\mathbf{a_F} = g_m\mathbf{h}$: $\mathbf{F}_t = t\mathbf{a_F} + \mathcal{O}(1)$. Therefore, temperature dynamics are asymptotically given by:

$$\dot{\mathbf{T}}_t = \mathbf{B}\mathbf{T}_t + t\mathbf{N}\mathbf{a_F} + \mathcal{O}(1).$$

Solving for $\mathbf{T}_t$:

$$\mathbf{T}_t = e^{t\mathbf{B}}\mathbf{T}_0 + \int_0^t e^{(t-s)\mathbf{B}} s\ \mathbf{N}\mathbf{a_F} ds + \mathcal{O}\left(\int_0^t e^{(t-s)\mathbf{B}} ds\right)$$

Recall that $\mathbf{B} = \mathbf{C}^{-1}(\mathbf{D} - \mathbf{L})$. The properties of $\mathbf{D}, \mathbf{L}$ imply that $\mathbf{B}$ is a Metzler matrix with $\mathbf{D}$ irreducible. Hence, $\mathbf{B}$ has eigenvalues with strictly negative real parts, implying that $e^{t\mathbf{B}}\mathbf{T}_0$ converges exponentially to zero. Similarly $\int_0^t e^{(t-s)\mathbf{B}} ds$ converges exponentially to a constant. Integrating by parts since $\mathbf{B}$ is invertible implies that: $\int_0^t e^{(t-s)\mathbf{B}} s\ \mathbf{N} ds = -t\mathbf{B}^{-1}\mathbf{N} + \mathcal{O}(1)$. Therefore:

$$\mathbf{T}_t = -tg_m\mathbf{B}^{-1}\mathbf{N}\mathbf{h} + \mathcal{O}(1) = tg_m(\mathbf{L} - \mathbf{D})^{-1}\mathbf{h} + \mathcal{O}(1) = tg_m\boldsymbol{\theta} + \mathcal{O}(1),$$

where the thermal capacity matrix cancels out in the $\mathbf{B}$ inverse and the definition of $\mathbf{N}$. Since $\mathbf{L} - \mathbf{D}$ is a nonsingular M-matrix, its inverse is a positive matrix, in that it maps the positive orthant into itself. Thus, $\boldsymbol{\theta} \geq 0$ with at least some strictly positive entries.

Claim: Y_t grows at rate $g_Y = g_\chi + g_\sigma + \frac{g-\omega}{1+\Delta}$. The condition $g_m > 0$ becomes $g - \omega > 0$.

To prove this claim, we turn to the asymptotic dynamics of Z_t and check that asymptotically linear temperature delivers constant growth. $\log Z_t = \log Z_0 + gt - \boldsymbol{\zeta}'\left(\mathbf{T}_t - \mathbf{T}_{eq}\right)$.

$$\log Z_t = \log Z_0 + gt - \Delta g_m t + \mathcal{O}(1),$$

where $\Delta = \boldsymbol{\zeta}'\boldsymbol{\theta}$.

Using the expression for g_Y and g_m, we obtain $g_Y = g - \Delta(g_Y - g_\chi - g_\sigma) - \pi g_\chi$. Rearranging: $(1+\Delta)g_Y = g - \omega + (1+\Delta)(g_\chi + g_\sigma)$, which delivers the expression in Proposition 1. The condition $g_m > 0$ then becomes $g - \omega > 0$.

Case 2: Falling emissions $g_m < 0$. The derivations in Case 2 mirror those from Case 1. Emissions are asymptotically equal to zero. By mass conservation, carbon masses in all sinks are asymptotically constant. Radiative forcing is therefore also asymptotically constant. With constant forcing, temperature is asymptotically constant, and so are damages. Hence, Z_t grows asymptotically at constant rate g. Using the expression for g_Y as a function of g_Z, the growth rate of output is simply: $g_Y = g - \pi g_X$. Using the expression for g_m as a function of g_Y, the condition that $g_m < 0$ is thus equivalent to: $g < \omega$.

Case 3: Constant emissions $g_m = 0$. The derivations in Case 3 mirror those from Case 1. Carbon grows asymptotically linearly in time since the dominant term in $\int_0^t e^{s\mathbf{A}} ds$ corresponds to the zero eigenvalue. As a consequence, radiative forcing grows log-linearly in time: $\mathbf{F}_t = \mathbf{h} \log t + \mathcal{O}(1)$. Similar calculations to those above then imply that temperatures grow asymptotically log-linearly: $\mathbf{T}_t = \boldsymbol{\theta} \log t + \mathcal{O}(1)$. This property implies that Z_t asymptotically grows at a constant rate g, and economic variables grow as in Case 1. $g_m = 0$ then corresponds to $g = \omega$.

A.5 Proof of Propositions 2 and 4

In this section we prove Proposition 2 under log utility, and we prove its variant under general CRRA utility, which we state below.

Away from logarithmic utility, the classification of efficient growth paths is richer than under log utility, although it retains the same overall structure. We provide a partial characterization below.

Proposition 4 (Efficient Growth With CRRA Utility). *If $\Delta \leq \pi$, then $\overline{g}_m = +\infty$ for any γ, and growth rates coincide with those in the laissez-faire from Proposition 1. If $\Delta > \pi$, there exists a threshold $\overline{\rho} > 0$ such that, if $\rho < \overline{\rho}$ and $\gamma \neq 1$:*

- *If $\gamma > 1$: the efficient allocation admits a unique asymptotic constant growth path. There exists a unique finite threshold $\overline{g}_m$ such that the results from Proposition 2 apply conditional on $\overline{g}_m$.*

- *If $\gamma < 1$: any asymptotic growth path admits the structure from Proposition 2(i) or (ii).*

For $\gamma = 1$, Proposition 2 applies for all $\rho > 0$.

We organize the proof in several steps.

A.5.1 Optimality conditions

The current-value Hamiltonian of the planner's problem is:

$$
\begin{aligned}
\mathcal{H} = U(C_t) + \lambda_t \Big(F\big(Z(\mathbf{T}_t), K_t, 1, E_t\big) - \chi_t E_t - C_t - \delta K_t \Big) \\
- \boldsymbol{\mu}_t' \Big(\mathbf{A}\mathbf{M}_t + \sigma_t E_t \mathbf{m} \Big) - \boldsymbol{\nu}_t' \Big(\mathbf{B}\mathbf{T}_t + \mathbf{N}\mathbf{F}_{eq} + \mathbf{N}\mathbf{H}\big(\log \mathbf{M}_t - \log \mathbf{M}_{eq} \big) \Big),
\end{aligned}
$$

where we normalize $\boldsymbol{\mu}_t, \boldsymbol{\nu}_t$ so that they are positive. The optimality condition and co-state equation for capital accumulation lead to the Euler equation: $\gamma \dot{C}_t / C_t = F_K(Z_t, K_t, 1, E_t) - \delta - \rho$ and $\lambda_t = u'(C_t)$. The optimality condition with respect to energy is:

$$
\chi_t = \psi \frac{Y_t}{E_t} - \frac{\sigma_t}{u'(C_t)} \boldsymbol{\mu}_t' \mathbf{m}.
$$

The co-state equations for carbon and temperature are:

$$
\rho \boldsymbol{\mu}_t - \dot{\boldsymbol{\mu}}_t = \mathbf{A}' \boldsymbol{\mu}_t + \mathcal{D}(\mathbf{M}_t)^{-1} \mathbf{P} \boldsymbol{\nu}_t \qquad\qquad \rho \boldsymbol{\nu}_t - \dot{\boldsymbol{\nu}}_t = u'(C_t) Y_t \overline{\boldsymbol{\zeta}} + \mathbf{B}' \boldsymbol{\nu}_t
$$

where $\mathcal{D}(\mathbf{M}) = \mathbf{diag}((\mathbf{M}_i)_i)$ denotes the diagonal matrix with the vector $\mathbf{M}$ on the diagonal, $\mathbf{P} = (\mathbf{N}\mathbf{H})'$, and $\overline{\boldsymbol{\zeta}} \equiv (1-\alpha)(1-\psi)\boldsymbol{\zeta}$. The transversality conditions are:

$$
e^{-\rho t} u'(C_t) K_t \to 0 \qquad\qquad e^{-\rho t} \boldsymbol{\mu}_t \circ \mathbf{M}_t \to 0 \qquad\qquad e^{-\rho t} \boldsymbol{\nu}_t \circ \mathbf{T}_t \to 0,
$$

where $\circ$ denotes the element-wise product.

A.5.2 Constant growth path

Economic module. The economic module behaves as in the laissez-faire, except for energy demand. Consider an asymptotic growth path. Let g_Z be the asymptotic growth rate of productivity. The resource constraint requires that the growth rates of output, consumption and capital are equal: $g_Y = g_C = g_K$. The structure of the production function then requires: $g_Y = \alpha(1-\psi)g_Y + (1-\alpha)(1-\psi)g_Z + \psi g_E$. The capital transversality condition implies: $(1-\gamma)g_Y < \rho$. However, energy demand now depends on whether the extraction cost or the shadow value of emissions grows faster in the long run. To determine it, we turn to the climate module.

Shadow values. We now turn to the asymptotic behavior of the planner's shadow values. We organize the discussion into a sequence of claims.

Claim: $\boldsymbol{\nu}_t = C_t^{-\gamma} Y_t \overline{\boldsymbol{\nu}}$, where $\overline{\boldsymbol{\nu}} = [\overline{\mathbf{B}} - (1-\gamma)g_Y \mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}$. $\boldsymbol{\nu}_t > 0$ if $\boldsymbol{\zeta} > 0$.

To prove this claim and accurately account for long-run stability, we solve the costate equation for $\boldsymbol{\nu}_t$ forward. To leading order:

$$
\begin{aligned}
\boldsymbol{\nu}_t &= e^{t\overline{\mathbf{B}}}\boldsymbol{\nu}_0 - C_0^{-\gamma} Y_0 e^{(1-\gamma)g_Y t}\int_0^t e^{s[\overline{\mathbf{B}}-(1-\gamma)g_Y\mathbf{I}]}\overline{\boldsymbol{\zeta}} ds \\
&= e^{t\overline{\mathbf{B}}}\boldsymbol{\nu}_0 - C_0^{-\gamma} Y_0 e^{(1-\gamma)g_Y t} e^{t[\overline{\mathbf{B}}-(1-\gamma)g_Y\mathbf{I}]}[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}} \\
&+ C_0^{-\gamma} Y_0 e^{(1-\gamma)g_Y t}[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}} \\
&= e^{t\overline{\mathbf{B}}}\left\{\boldsymbol{\nu}_0 - C_0^{-\gamma} Y_0[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}\right\} + C_t^{-\gamma} Y_t[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}
\end{aligned}
$$

where $\overline{\mathbf{B}} = \rho\mathbf{I} - \mathbf{B}'$. $\overline{\mathbf{B}}$ has eigenvalues with strictly positive real part. From the capital transversality condition, so does $\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}$. Therefore, the first component grows at rate ρ unless $\boldsymbol{\nu}_0$ is such that the bracket is equal to zero. Since temperature is weakly increasing over time, growth in the first component at rate ρ violates the temperature transversality condition. Therefore, $\boldsymbol{\nu}_0$ is such that the first component is zero, and to leading order: $\boldsymbol{\nu}_t = C_t^{-\gamma} Y_t \overline{\boldsymbol{\nu}}$, where $\overline{\boldsymbol{\nu}} = [\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}$. From the capital transversality condition, $\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}$ is a non-singular M-matrix. Therefore, its inverse is a positive matrix, so that $\overline{\boldsymbol{\nu}} > 0$, and hence $\boldsymbol{\nu}_t > 0$, if $\boldsymbol{\zeta} > 0$.

Claim: $\boldsymbol{\mu}_t = \frac{Y_t}{C_t^{\gamma} M_t}\overline{\boldsymbol{\mu}}$, where $\overline{\boldsymbol{\mu}} = [\overline{\mathbf{A}} - (1-\gamma)g_Y\mathbf{I} + g_M\mathbf{I}]^{-1}\mathcal{D}(\mathbf{M}_0/M_0)^{-1}\mathbf{P}[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}$, where g_M denotes the exponential growth rate of carbon stocks, and M_t denotes an arbitrary but fixed entry of the carbon sink vector.[12]

To prove this claim, substitute the expression for $\boldsymbol{\nu}_t$ into the equation for $\boldsymbol{\mu}_t$ and solve

[12]For now we abstract from the knife-edge condition under which emissions are exactly constant and therefore carbon stocks grow linearly.

forward to leading order:

$$\begin{aligned}\boldsymbol{\mu}_t &= e^{t\overline{\mathbf{A}}}\boldsymbol{\mu}_0 - C_0^{-\gamma}Y_0 e^{[(1-\gamma)g_Y - g_M]t}\int_0^t e^{s[\overline{\mathbf{A}}-(1-\gamma)g_Y+g_M]}\cdot\mathcal{D}(\mathbf{M}_0)^{-1}\mathbf{P}\overline{\boldsymbol{\nu}}ds\\ &= e^{t\overline{\mathbf{A}}}\boldsymbol{\mu}_0 - C_0^{-\gamma}Y_0 e^{[(1-\gamma)g_Y - g_M]t}e^{t[\overline{\mathbf{A}}-(1-\gamma)g_Y+g_M]}[\overline{\mathbf{A}} - g_M\mathbf{I} + (1-\gamma)g_Y\mathbf{I}]^{-1}\mathcal{D}(\mathbf{M}_0)^{-1}\mathbf{P}\overline{\boldsymbol{\nu}}\\ &+ C_0^{-\gamma}Y_0 e^{[(1-\gamma)g_Y - g_M]t}[\overline{\mathbf{A}} - (1-\gamma)g_Y\mathbf{I} + g_M\mathbf{I}]^{-1}\mathcal{D}(\mathbf{M}_0)^{-1}\mathbf{P}\overline{\boldsymbol{\nu}}\end{aligned}$$

where $\overline{\mathbf{A}} = \rho\mathbf{I} - \mathbf{A}'$. Since $\overline{\mathbf{A}}$ has eigenvalues with strictly positive real part, transversality imposes that the sum of the first two components is zero. Therefore: $\boldsymbol{\mu}_t = \frac{Y_t}{C_t^\gamma M_t}\overline{\boldsymbol{\mu}}$, where $\overline{\boldsymbol{\mu}} = [\overline{\mathbf{A}} - (1-\gamma)g_Y\mathbf{I} + g_M\mathbf{I}]^{-1}\mathcal{D}(\mathbf{M}_0/M_0)^{-1}\mathbf{P}[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}$, and M_t denotes an arbitrary but fixed entry of the carbon sink vector.

Energy use. We are now ready to return to energy use, which satisfies to leading order:

$$\left(\psi - \phi\frac{E_t^m}{M_t}\right)\frac{Y_t}{\chi_t E_t} = 1, \tag{15}$$

where $\phi = \overline{\boldsymbol{\mu}}'\mathbf{m}$. There are two possibilities: either $\psi > \phi\frac{E_t^m}{M_t}$, or $\psi = \phi\frac{E_t^m}{M_t}$. $\psi < \phi\frac{E_t^m}{M_t}$ would lead to a contradiction.

We organize the remainder of the proof into three cases, corresponding to whether emissions are falling, constant or growing. We characterize which parameters imply that each case arises as part of the proof.

A.5.3 Falling emissions $g_m < 0$

If $g_m < 0$, carbon stocks are asymptotically constant over time. E_t^m/M_t converges to zero. Temperature and damages are asymptotically constant, so that $g_Z = g$. Energy use then mirrors the laissez-faire, with $g_Y = g_\chi + g_E$ and so $g_Y = g - \pi g_\chi$. The condition $g_m < 0$ becomes $g < \omega$.

A.5.4 Constant emissions $g_m = 0$.

Asymptotically, carbon stocks grow linearly over time. Temperature and damages grow log-linearly, so that $g_Z = g$. E_t^m/M_t converges to zero. As in the laissez-faire, we obtain $g_Y = g - \pi g_\chi$. The condition $g_m = 0$ becomes $g = \omega$.

A.5.5 Growing emissions $g_m > 0$.

We organize the remaining proof in a sequence of claims.

Claim: Asymptotically, $g_M = g_m$ and $\mathbf{M}_t = E_t^m[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}$.

To prove this claim, solve carbon stocks forward to leading order:

$$\begin{aligned}\mathbf{M}_t &= e^{t\mathbf{A}}\mathbf{M}_0 + E_0^m e^{tg_m}\int_0^t e^{-s[g_m\mathbf{I}-\mathbf{A}]}\mathbf{m}ds\\ &= e^{t\mathbf{A}}\mathbf{M}_0 - E_0^m e^{t\mathbf{A}}[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m} + E_t^m[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}.\end{aligned}$$

Because $g_m > 0$ and the eigenvalues of $\mathbf{A}$ have weakly negative real parts, the third component dominates the other two. Hence, asymptotically, $g_M = g_m$, and: $\mathbf{M}_t = E_t^m[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}$.

Claim: $\sigma_t\tau_t = \overline{\boldsymbol{\zeta}}'[\eta\mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H}\mathcal{D}\left([g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right)^{-1}[\eta\mathbf{I} + g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\ \frac{Y_t}{E_t}$.

To prove this claim, let $\mathbf{e}$ be the vector that selects M_t out of $\mathbf{M}_t$. Then: $M_t = \mathbf{e}'\mathbf{M}_t = E_t^m\mathbf{e}'[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}$, and so to leading order:

$$\frac{E_t^m}{M_t} = \frac{1}{\mathbf{e}'[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}} \qquad\qquad \frac{\mathbf{M}_0}{M_0} = \frac{[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}}{\mathbf{e}'[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}}.$$

Recall that $\overline{\boldsymbol{\mu}} = [\overline{\mathbf{A}} - (1-\gamma)g_Y\mathbf{I} + g_m\mathbf{I}]^{-1}\mathcal{D}(\mathbf{M}_0/M_0)^{-1}\mathbf{P}[\overline{\mathbf{B}} - (1-\gamma)g_Y\mathbf{I}]^{-1}\overline{\boldsymbol{\zeta}}$. Since $g_m > 0$ and $\rho > (1-\gamma)g_Y$, all the matrices are non-singular irreducible M-matrices and therefore have positive inverses. Then $\overline{\boldsymbol{\mu}} > 0$ if $\boldsymbol{\zeta} > 0$. We obtain:

$$\begin{aligned}\phi \times \frac{E_t^m}{M_t} &= \left\{\overline{\boldsymbol{\zeta}}'[(\rho - (1-\gamma)g_Y)\mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H}\mathcal{D}\left(\frac{[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}}{\mathbf{e}'[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}}\right)^{-1}[(\rho - (1-\gamma)g_Y + g_m)\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right\}\\ &\quad\times\frac{1}{\mathbf{e}'[g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}}\\ &= \overline{\boldsymbol{\zeta}}'[\eta\mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H}\mathcal{D}\left([g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right)^{-1}[\eta\mathbf{I} + g_m\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\end{aligned}$$

where in the second equality, the second term simplifies with $\mathcal{D}$, and we denoted $\eta =$

$\rho - (1-\gamma)g_Y > 0$.[13] We have thus shown that:

$$\sigma_t \tau_t = \overline{\boldsymbol{\zeta}}'[\eta \mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H}\mathcal{D}\left([g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right)^{-1}[\eta \mathbf{I} + g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\,\frac{Y_t}{E_t}. \tag{16}$$

We are now ready to consider the main cases of Propositions 2 and 4. We start with Proposition 2, which corresponds to log utility.

A.5.6 Log Utility $\gamma = 1$

We impose log utility $\gamma = 1$ in this section.

No-cap regime: $\psi > \phi E_t^m / M_t$. In that case, energy use mirrors the laissez-faire (third case of Proposition 1), and we obtain the same growth rates. The condition $g_m > 0$ becomes $g > \omega$. However, we will have to verify which parameters lead to $\psi > \phi E_t^m / M_t$.

Cap regime: $\psi = \phi E_t^m / M_t$. Then:

$$\begin{aligned}
\frac{\phi}{\psi}\frac{E_t^m}{M_t} &= \overline{\boldsymbol{\zeta}}'[\rho \mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H}\mathcal{D}\left([g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right)^{-1}[\rho \mathbf{I} + g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m} \\
&= \frac{\overline{\boldsymbol{\Delta}}}{\pi}\mathcal{D}\left([g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}\right)^{-1}[\rho \mathbf{I} + g_m \mathbf{I} - \mathbf{A}]^{-1}\mathbf{m} \\
&\equiv f(g_m).
\end{aligned}$$

Denote $\boldsymbol{\xi}' = \frac{\overline{\boldsymbol{\Delta}}}{\pi} = \frac{\overline{\zeta}'}{\psi}[\rho \mathbf{I} - \mathbf{B}]^{-1}\mathbf{N}\mathbf{H} \geq 0$, with some inequalities strict. We now leverage the following result:

Lemma 1. *f is continuous, strictly increasing, $f(0) = 0$ and $f(+\infty) = \sum_i \xi_i$.*

Proof. See Online Appendix E.1. □

When $\sum_i \xi_i > 1$, there is a unique value $\overline{g}_m$ such that $f(\overline{g}_m) = 1$. This condition immediately rewrites as: $\sum_i \overline{\Delta}_i > \pi$. This value $\overline{g}_m$ only depends on parameters. The structure and properties of f immediately imply that $\overline{g}_m$ is decreasing in each component ξ_i and is independent of g. When $\sum_i \xi_i < 1$, we define by convention: $\overline{g}_m = +\infty$. We always have

[13] We can also re-express this formula. Let $\widetilde{\mathbf{A}} = g_m \mathbf{I} - \mathbf{A}$ and $a = \widetilde{\mathbf{A}}^{-1}\mathbf{m}$. The last two components are equal to:

$$\mathcal{D}(a)^{-1}[\widetilde{\mathbf{A}} + \eta \mathbf{I}]^{-1}\widetilde{\mathbf{A}}a = \mathbf{I} - \eta\mathcal{D}(a)^{-1}(\widetilde{\mathbf{A}} + \eta \mathbf{I})^{-1}a$$

where we have used the resolvent identity: $[\widetilde{\mathbf{A}} + \eta \mathbf{I}]^{-1}\widetilde{\mathbf{A}} = \mathbf{I} - \eta(\widetilde{\mathbf{A}} + \eta \mathbf{I})^{-1}$.

$g_m < \overline{g}_m$ in the no-cap regime, and $g_m = \overline{g}_m$ in the cap regime.

We are now in a position to characterize which parameters determine when the cap binds. When the cap does not bind, we have $g_m = \frac{g-\omega}{1+\Delta}$ as in the laissez-faire. Therefore, the cap does not bind when $g < \omega + (1+\Delta)\overline{g}_m$ under log utility.

The cap binds when $g \geq \omega + (1+\Delta)\overline{g}_m$. We can then solve for $g_E = \overline{g}_m + g_\sigma$. From the production function: $g_Y = \alpha(1-\psi)g_Y + (1-\alpha)(1-\psi)g_Z + \psi g_E$, and so $[1-\alpha(1-\psi)]g_Y = (1-\alpha)(1-\psi)g_Z + \psi(\overline{g}_m + g_\sigma)$. Noting that $1-\alpha(1-\psi) = (1-\alpha)(1-\psi)(1+\pi)$, we obtain: $(1+\pi)g_Y = g_Z + \pi g_E = g - \Delta\overline{g}_m + \pi(\overline{g}_m + g_\sigma)$. Re-arranging, we obtain:

$$g_Y = \frac{g + (\pi - \Delta)\overline{g}_m + \pi g_\sigma}{1+\pi}.$$

A.5.7 CRRA Utility $\gamma \neq 1$

In this section, we consider the case of general CRRA utility with $\gamma \neq 1$.

No-cap regime: $\psi > \phi E_t^m / M_t$. As with log utility, energy use mirrors the laissez-faire (third case of Proposition 1), and we obtain the same growth rates. The condition $g_m > 0$ becomes $g > \omega$. However, we will have to verify which parameters lead to $\psi > \phi E_t^m / M_t$.

Cap regime: $\psi = \phi E_t^m / M_t$. Define: $z(x) = \overline{\boldsymbol{\zeta}}'[x\mathbf{I} - \mathbf{B}]^{-1}$. Denote by $N_i \geq 0$ the diagonal elements of $\mathbf{NH}$, with some inequalities strict. Then:

$$\overline{f}(\eta, x) = \sum_{i=1}^{\mathcal{T}} z_i(\eta) N_i \frac{y_i(\eta + x)}{y_i(x)}.$$

By the same arguments as for Lemma 1, $\overline{f}$ is increasing in x, decreasing in η. Given η, $\overline{f}$ limits to 0 and $\sum_i z_i(\eta)N_i$ when x varies between 0 (its lower bound because of the capital transversality condition) and $+\infty$. Note that $\sum_i z_i(\eta)N_i = \overline{\boldsymbol{\zeta}}'[\eta\mathbf{I} - \mathbf{B}]^{-1}\mathbf{Nh} = (1-\alpha)(1-\psi)\,\boldsymbol{\zeta}'(\eta\mathbf{C} + \mathbf{L} - \mathbf{D})^{-1}\mathbf{h}$: it is decreasing in η and equals $(1-\alpha)(1-\psi)\Delta$ at $\eta = 0$.

Therefore, there exists a unique strictly increasing $X(\eta)$ such that $\overline{f}(\eta, X(\eta)) = \psi$, which we normalize to be infinite if $\psi \geq \sum_i z_i(\eta)N_i$. Since $\overline{f} < \sum_i z_i(\eta)N_i \leq (1-\alpha)(1-\psi)\Delta$ for all (η, x), a finite $X(\eta)$ for some $\eta \geq 0$ requires $(1-\alpha)(1-\psi)\Delta > \psi$, that is, $\Delta > \pi$.

If $\Delta \leq \pi$, then $X(\eta) = +\infty$ for every η: the equation $\psi = \phi E_t^m / M_t$ has no solution, the capped regime never arises, and growth rates coincide with the laissez-faire for any

γ, with $\overline{g}_m = +\infty$, as stated in Proposition 4. We therefore assume $\Delta > \pi$ in what follows.

Under $\Delta > \pi$, there exists $\overline{\eta} > 0$ such that $\sum_i N_i z_i(\eta) > \psi$ for $\eta < \overline{\eta}$. We also have $X(0) = 0$: as $\eta \to 0$, the zero eigenvalue of $\mathbf{A}$ implies $y_i(u) \sim v_i^\infty/u$ near $u = 0$, where $v^\infty > 0$ denotes the stationary distribution of $\mathbf{A}$, so that $y_i(\eta + x)/y_i(x) \to x/(\eta + x)$ and solving $\overline{f} = \psi$ yields $x \approx \eta\pi/(\Delta - \pi) \to 0$. Hence X increases from 0 to $+\infty$ for $\eta \in [0, \overline{\eta})$. To avoid a lengthier discussion of cases below, we therefore impose $\rho < \overline{\eta}$ in what follows. $\overline{\eta}$ is our threshold $\overline{\rho}$.[14]

Substituting into the output growth expression, we obtain:

$$(1+\pi)g_Y = g + (\pi - \Delta)X(\rho - (1-\gamma)g_Y) + \pi g_\sigma.$$

If $\boldsymbol{\gamma > 1}$: The right-hand side is decreasing in g_Y and spans $g + \pi g_\sigma > 0$ to $-\infty$ as g_Y spans $-\frac{\rho}{\gamma-1} < 0$ to $\frac{\overline{\eta}-\rho}{\gamma-1} > 0$, its lowest and largest possible values. Therefore, there is always a unique solution $g_Y > 0$, with an associated finite threshold $\overline{g}_m$.

If $\boldsymbol{\gamma < 1}$: The right-hand side is increasing in g_Y and spans $-\infty$ to $g + \pi g_\sigma$ as g_Y spans $\frac{\rho-\overline{\eta}}{1-\gamma} < 0$ to $\frac{\rho}{1-\gamma} > 0$, its lowest and largest possible values. Hence, there can be no solution, exactly one solution in the knife-edge tangent case, or two or more solutions depending on the exact shape of X.

In both cases, the implicit solution g_Y always corresponds to a $\overline{g}_m$ such that $g_Y = \frac{g+(\pi-\Delta)\overline{g}_m+\pi g_\sigma}{1+\pi}$, because this last expression only uses the production function and the definition of emissions.

A.5.8 Comparison of growth rates

We have in the high emissions regime, using the expression for g_Y as a function of g_m in the laissez-faire:

$$g_Y^{\mathrm{SP}} > g_Y^{\mathrm{LF}} \Leftrightarrow \frac{g + (\pi - \Delta)\overline{g}_m + \pi g_\sigma}{1+\pi} > \frac{g + (\pi - \Delta)g_m + \pi g_\sigma}{1+\pi} \Leftrightarrow (\Delta - \pi)(g_m - \overline{g}_m) > 0$$

[14] This restriction is not binding in practice. To see why, in a one-dimensional climate system, $\sum_i N_i z_i(\eta) = \frac{(1-\alpha)(1-\psi)\zeta n h}{\eta + n\ell}$, and so $\overline{\eta} = \frac{\zeta n h}{\pi} - n\ell$, which is positive precisely when $\Delta = \zeta h/\ell > \pi$. We consider damages between 0.036 and 0.25. π is of order 0.075. $n = 0.14$, $h \approx 5$ and $\ell \approx 1$. Therefore, $\overline{\eta}$ is of order 0.2 to 2.

where $g_m = \frac{g-\omega}{1+\Delta} > \overline{g}_m$. Now, recall that by the properties of M-matrices, $\overline{\Delta} < \Delta$. $\overline{g}_m$ is finite if $\overline{\Delta} > \pi$, and so $\Delta > \pi$. Therefore, $g_Y^{\mathrm{SP}} > g_Y^{\mathrm{LF}}$.

A.6 Proof of Proposition 3

We consider a possibly nonlinear damage function $D(T_t)$ such that: $\log Z_t = \log Z_0 + gt - D(T_t)$. For brevity, we consider a unidimensional climate system, but our proof extends immediately to a multidimensional climate system.

A.6.1 Laissez-faire

The asymptotic behavior of temperature conditional on emissions is independent of the damage function. We thus recall:

- If $g_m < 0$, temperature T_t stabilizes.
- If $g_m = 0$, temperature grows log-linearly: $T_t = \theta \log t + \mathcal{O}(1)$
- If $g_m > 0$, temperature grows linearly: $T_t = t g_m \theta + \mathcal{O}(1)$.

We start from the accounting identities that hold away from a growth path: $g_Y = g_Z - \pi g_\chi$ and $g_Y = g_m + g_\chi + g_\sigma$, where we omit time subscripts on growth rates when unambiguous. Together, they imply: $g_m = g_Z - \omega$. In addition, we have: $g_Z = g - D'(T_t)\dot{T}_t$.

Case 1. $D'(T) \to 0$ as $T \to \infty$. If $g_m < 0$, $g_Z = g$. If $g_m = 0$, $|g_Z - g| = |D'(T_t)|\theta/t \to 0$. If $g_m > 0$, $|g_Z - g| = |D'(T_t)|\theta g_m \to 0$. In all cases, technological growth dominates in the long run. In that case, climate damages have no long-run effect on growth rates, and the expressions from Proposition 1(i) apply regardless of the comparison between g and ω.

Case 2. $D'(T) \to +\infty$ as $T \to \infty$. We start with a few restrictions when $g > \omega$. Start from $g_m = g - \omega - D'(T_t)\dot{T}_t$. Then T_t must go to infinity. Indeed, otherwise if T_t is bounded above (and long-run dynamics are monotone), $\dot{T}_t = 0$ and $D'(T_t)$ converges to a constant, so that $g_m = g - \omega > 0$, a contradiction.

Case 2.1. $g_m > 0$. Then: $g_Z = g - D'(T_t)\theta g_m \to -\infty$. Hence $g_Y \to -\infty$, and so $g_m \to -\infty$, a contradiction. Thus, this case can never arise.

In addition, if emissions growth is not necessarily exponential but bounded below by a positive constant $g_{mt} > 0$, then we obtain the same contradiction as just above. Hence, emissions growth must be weakly negative (including zero) asymptotically. However, there can be positive growth at finite t, as long as the growth rate converges to zero. By the same token, we cannot have $g_{Mt} > 0$ bounded below, for the carbon stock.

Case 2.2. $g_m < 0$. Temperatures stabilize. Therefore: $g_Z = g - D'(T_t)\dot{T}_t$. Hence $g_Z = g$ in the long run, and so $g_m = g - \omega$. If $g - \omega < 0$, the equilibrium is compatible with Proposition 1(i), which therefore is the equilibrium. If $g > \omega$, we obtain a contradiction. Hence, $g > \omega$ is incompatible with $g_m < 0$ and with $g_m > 0$.

Thus, we now focus on the parametric case $g > \omega$. We must have strictly increasing temperature T_t that does not converge. We allow the emissions growth rate g_{mt} to depend on time. In the unidimensional climate system: $F_t = \overline{F} + h \log M_t$ and $\dot{T}_t = nF_t - \overline{\ell} T_t$. Therefore:

$$\dot{T}_t = n\overline{F} + nh \log M_t - \overline{\ell} T_t.$$

Since we operate under the case $\dot{M}_t / M_t = g_{Mt} \to 0$ from above, we guess and verify that, to leading order, the solution is:

$$T_t = \frac{nh}{\overline{\ell}} \log M_t \equiv \frac{1}{a} \log M_t,$$

where $a = \frac{\overline{\ell}}{nh} = \frac{\ell}{h}$. This condition is indeed compatible with $\dot{T}_t \to 0$. Further define $b = g - \omega > 0$. Then, to leading order: $a\dot{T}_t = \frac{\dot{M}_t}{M_t} = g_{mt}$. We thus have:

$$a\dot{T}_t = g_{mt} = b - D'(T_t)\dot{T}_t, \qquad \text{and so} \qquad (a + D'(T_t))\dot{T}_t = b.$$

Hence, to leading order:

$$aT_t + D(T_t) = bt.$$

Let $\overline{T}_t$ be such that $D(\overline{T}_t) = bt$. And let $T_t = \overline{T}_t + \widehat{T}_t$, with $\widehat{T}_t \ll \overline{T}_t$. Expanding the T_t

equation, we obtain to leading order: $a\overline{T}_t + a\widehat{T}_t + D'(\overline{T}_t)\widehat{T}_t = 0$, and so:

$$\widehat{T}_t = -\frac{a\overline{T}_t}{a + D'(\overline{T}_t)} \qquad\qquad T_t = \overline{T}_t - \frac{a\overline{T}_t}{a + D'(\overline{T}_t)}.$$

We then obtain, to leading order:

$$g_{mt} = a\dot{T}_t = \frac{ab}{a + D'(T_t)} = \frac{ab}{a + D'(\overline{T}_t)} \to 0.$$

Thus, to leading order:

$$g_{Yt} = g_\chi + g_\sigma + g_{mt} = g_\chi + g_\sigma + \frac{ab}{a + D'(\overline{T}_t)} \to g_\chi + g_\sigma.$$

When $D(T) = \zeta T^{1+d}$, we obtain: $\overline{T}_t = \left(\frac{bt}{\zeta}\right)^{\frac{1}{1+d}}$. When $D(T) = \zeta e^{dT}$, we obtain to leading order: $\overline{T}_t = \frac{1}{d}\log t$.

A.6.2 Efficient Allocation

We start from the accounting identity that holds away from a growth path: $g_Y = g_Z - \pi g_\chi$.

Case 1. $D'(T) \to 0$ as $T \to \infty$. If $g_m < 0$, $g_Z = g$. If $g_m = 0$, $|g_Z - g| = |D'(T_t)|\theta/t \to 0$. If $g_m > 0$, $|g_Z - g| = |D'(T_t)|\theta g_m \to 0$. In all cases, technological growth dominates in the long run. In that case, climate damages have no long-run effect on growth rates, and the expressions from the low emissions regime from Propositions 2 and 4 apply for all values of g, g_χ, g_σ.

Case 2. $D'(T) \to +\infty$ as $T \to \infty$.

Social cost of carbon with growing emissions. We must now characterize the social cost of carbon. Denote $\overline{\zeta}_t = (1-\alpha)(1-\psi)D'(T_t)$, which is weakly rising over time. Our calculation holds regardless of the sign of g_m.

Then:

$$(\rho + \overline{\ell} - g_\nu)\nu_t = u'(C_t)Y_t\overline{\zeta}_t \qquad\qquad (\rho - g_\mu)\mu_t = \frac{nh\nu_t}{M_t} = \frac{nhY_t\overline{\zeta}_t}{(\rho + \overline{\ell} - g_\nu)C_t^\gamma M_t},$$

where we normalize μ_t, ν_t so that they are positive, as in Appendix A.5. We then obtain: $g_\nu = (1-\gamma)g_Y + g_\zeta$ and $g_\mu = (1-\gamma)g_Y + g_\zeta - g_M$, and so:

$$\mu_t = \frac{nh}{(\rho + \bar{\ell} - (1-\gamma)g_Y - g_\zeta)(\rho - (1-\gamma)g_Y - g_\zeta + g_M)} \frac{Y_t \bar{\zeta}_t}{C_t^\gamma M_t} \equiv \frac{nh}{\overline{M}(g_Y, g_M, g_\zeta)} \frac{Y_t \bar{\zeta}_t}{C_t^\gamma M_t}$$

Recall the energy FOC: $\psi = \frac{\chi_t E_t}{Y_t} + \frac{E_t^m C_t^\gamma \mu_t}{Y_t}$. Then:

$$\frac{E_t^m C_t^\gamma \mu_t}{Y_t} = \frac{nh}{\overline{M}(g_Y, g_M, g_\zeta)} \bar{\zeta}_t \frac{E_t^m}{M_t} = \frac{nh \max\{g_m, 0\}}{\overline{M}(g_Y, g_m, g_\zeta)} \bar{\zeta}_t,$$

where the second equality uses the relationship between E_t^m and M_t when $g_m > 0$. Hence, the energy FOC becomes:

$$\left\{ \psi - \frac{nh \max\{g_m, 0\}}{\overline{M}(g_Y, g_m, g_\zeta)} \bar{\zeta}_t \right\} \frac{Y_t}{\chi_t E_t} = 1.$$

We denote $\Psi_t = \frac{nh \max\{g_m, 0\}}{\overline{M}(g_Y, g_m, g_\zeta)} \bar{\zeta}_t$.

Case 2.1. $g_m > 0$. Then: $g_Z = g - D'(T_t)\theta g_m \to -\infty$. Hence $g_Y \to -\infty$.

If $\psi > \Psi_t$ at all times, then $g_m = g_Y - g_\chi - g_\sigma \to -\infty$ and we obtain a contradiction. If $\psi = \Psi_t$ after some time, then $g_m = \bar{g}_m$ and $g_Y = \bar{g}_m + g_\chi + g_\sigma$ is constant, a contradiction. Therefore, this case can never arise.

Case 2.2. $g_m < 0$. Temperatures stabilize. Therefore: $g_Z = g - D'(T_t)\dot{T}_t$. Hence $g_Z = g$ in the long run. As before, we are in the low emissions regime, and so $g_m = g - \omega$. If $g - \omega < 0$, the equilibrium is compatible with the expressions from the low emissions regime from Propositions 2 and 4. If $g > \omega$, we obtain a contradiction. Hence, $g > \omega$ is incompatible with $g_m < 0$ and with $g_m > 0$.

We now focus on $g > \omega$. If $\psi > \Psi_t$ at all large times, the formulas from the laissez-faire apply and $g_{mt} \to 0$, $g_{Yt} \to g_\chi + g_\sigma$, and T_t grows to infinity. We must verify whether these conditions are compatible with the long-run behavior of Ψ_t. We have to leading order:

$$\begin{aligned}
\Psi_t &= \frac{nh}{\overline{M}(g_{Yt}, g_{mt}, g_{\zeta t})} g_{mt}\overline{\zeta}_t \\
&= \frac{\psi}{\pi} \frac{nh}{\overline{M}(g_\chi + g_\sigma, 0, g_{\zeta t})} \frac{abD'(\overline{T}_t)}{a + D'(\overline{T}_t)} \\
&= \frac{\psi}{\pi} \frac{anh}{\overline{M}(g_\chi + g_\sigma, 0, g_{\zeta t})} (g - \omega) \\
&= \frac{\psi}{\pi} \frac{n\ell}{\overline{M}(g_\chi + g_\sigma, 0, 0)} (g - \omega)
\end{aligned}$$

where the second line uses $\overline{\zeta}_t = (1-\alpha)(1-\psi)D'(\overline{T}_t) = \frac{\psi}{\pi} D'(\overline{T}_t)$, and with $g_{\zeta t} = \frac{D''(\overline{T}_t)}{D'(\overline{T}_t)} \dot{\overline{T}}_t = \frac{D''(\overline{T}_t)}{D'(\overline{T}_t)} \frac{b}{a + D'(\overline{T}_t)} = \frac{bD''(\overline{T}_t)}{\left(D'(\overline{T}_t)\right)^2} \to 0$, since, for damage functions that remain finite at all T we must have $D''(T)/(D'(T))^2 \to 0$.[15]

Evaluated at the laissez-faire growth rates, $\overline{M}(g_\chi + g_\sigma, 0, 0) = (\rho + \overline{\ell} + (\gamma - 1)(g_\chi + g_\sigma))(\rho + (\gamma-1)(g_\chi + g_\sigma))$. Then:

$$\Psi_t \to \overline{\Psi} \equiv \frac{\psi}{\pi} \frac{n\ell(g - \omega)}{\overline{M}(g_\chi + g_\sigma, 0, 0)}.$$

So if $\psi > \overline{\Psi}$, we are in the laissez-faire regime. This condition becomes:

$$g - \omega < \frac{\pi \overline{M}(g_\chi + g_\sigma, 0, 0)}{n\ell} = \frac{\pi\rho(\rho + \ell n)}{n\ell} = \vartheta,$$

where the second equality uses $\gamma = 1$ and matches the threshold ϑ in Proposition 3.

If $\psi < \overline{\Psi}$, then the planner chooses g_{mt} so that $\psi = \Psi_t$ at large times. For brevity, we characterize it when $\gamma = 1$ so that $\overline{M}$ does not depend on g_Y.

We have $a\dot{T}_t = g_{mt}$. Thus, we have:

$$\pi \overline{M}\left(a\dot{T}_t, \frac{D''(T_t)}{D'(T_t)} \dot{T}_t\right) = anh\dot{T}_t D'(T_t),$$

[15] If the limit is positive, solve the associated differential equation to obtain that D', and hence D, diverges on a finite support.

using again $\overline{\zeta}_t = \frac{\psi}{\pi} D'(T_t)$. Therefore:

$$\frac{anh}{\pi} D'(T_t)\dot{T}_t = \left(\widetilde{\rho} - \frac{D''(T_t)}{D'(T_t)}\dot{T}_t\right)\left(\rho - \left[\frac{D''(T_t)}{D'(T_t)} - a\right]\dot{T}_t\right).$$

with $\widetilde{\rho} = \rho + \overline{\ell}$, and $anh = \overline{\ell}$. We guess and verify that $\frac{D''(T_t)}{D'(T_t)}\dot{T}_t \to 0$. Under this guess, we have to leading order:

$$\left(\frac{\overline{\ell}}{\pi} D'(T_t) - \widetilde{\rho}a\right)\dot{T}_t = \rho\widetilde{\rho}.$$

Thus, to leading order:

$$D'(T_t)\dot{T}_t = \frac{\pi\rho\widetilde{\rho}}{\overline{\ell}} = \frac{\pi\rho(\rho + \overline{\ell})}{\overline{\ell}} = \vartheta.$$

This result then implies: $\frac{D''(T_t)}{D'(T_t)}\dot{T}_t \sim \frac{D''(T_t)}{(D'(T_t))^2} \to 0$ as T_t grows, and the guess is verified. Thus, to leading order we obtain:

$$g_{mt} = a\dot{T}_t = \frac{\pi\rho(\rho + \overline{\ell})}{nh}\frac{1}{D'(T_t)} \to 0.$$

From the production function and the definition of emissions, we obtain:

$$\begin{aligned}(1+\pi)g_{Yt} &= g - D'(T_t)\dot{T}_t + \pi(g_{mt} + g_\sigma) \\ &= g - \vartheta + \pi(g_{mt} + g_\sigma) \\ &\to g - \vartheta + \pi g_\sigma.\end{aligned}$$

Under our case $g > \omega$, we obtain in the long run:

$$\begin{aligned}(1+\pi)g_Y &= g - \omega + (1+\pi)g_\chi + g_\sigma - \vartheta + \pi g_\sigma \\ &= g - \omega + (1+\pi)(g_\chi + g_\sigma) - \vartheta \\ &= g - \omega - \vartheta + (1+\pi)g_Y^{\mathrm{LF}}.\end{aligned}$$

Therefore:

$$g_Y^{\mathrm{SP}} = g_Y^{\mathrm{LF}} + \frac{g - \omega - \vartheta}{1 + \pi}.$$

In our calibration, $\vartheta = \frac{\pi\rho(\rho+\bar{\ell})}{\bar{\ell}} = 0.001$. Hence, as soon as $g > \omega + \vartheta$, the planner achieves a strictly higher output growth rate than the laissez-faire.

B Climate Module

We parametrize the carbon cycle matrix $\mathbf{A}$ as in Folini et al. (2024):

$$\mathbf{A} = \begin{pmatrix} -a_1 & r_1 a_1 & 0 \\ a_1 & -(r_1 a_1 + a_2) & r_2 a_2 \\ 0 & a_2 & -r_2 a_2 \end{pmatrix}. \tag{17}$$

Quantitatively, we use the calibration from Folini et al. (2024): $a_1 = 0.054$, $a_2 = 0.0082$, $r_1 = \dfrac{M_{eq}^A}{M_{eq}^U}$, $r_2 = \dfrac{M_{eq}^U}{M_{eq}^L}$ and $M_{eq} = (607, 489, 1281)$, where M_{eq} is the pre-industrial equilibrium before any anthropogenic carbon emissions.

We follow Folini et al. (2024) for the calibration of the matrices $\mathbf{B}$ and $\mathbf{N}$:

$$\mathbf{B} = \begin{pmatrix} -nb_1 - n\ell & nb_1 \\ b_2 & -b_2 \end{pmatrix} \quad \mathbf{N} = \begin{pmatrix} n & 0 \\ 0 & 1 \end{pmatrix},$$

with the following calibration: $n = 0.137$, $b_1 = 0.73$, $b_2 = 0.00689$, $\ell = 1.06$, $F_{2XCO2} = 3.45$.

Equivalently, define T_{eq} as the equilibrium temperature vector when $F_t = 0$, and $\widehat{T}_t^X = T_t^X - T_{eq}^X$. The dynamics of the deviations from pre-industrial temperatures are given by:

$$\dot{\widehat{T}}_t^U = n \cdot \Big(-\ell \widehat{T}_t^U + b_1(\widehat{T}_t^L - \widehat{T}_t^U) + F_t \Big) \qquad \dot{\widehat{T}}_t^L = b_2(\widehat{T}_t^U - \widehat{T}_t^L).$$

Since we specify the model in levels rather than in deviations from steady-state, we also need to calibrate the constant forcing terms in the two layers by targeting temperature levels in the two layers (in Kelvin degrees). Specifically, we choose F_{eq}^U and F_{eq}^L to target

$T^U_{eq} = 288.15K$ (15C) and $T^L_{eq} = 278.15K$ (5C), by solving:

$$F^U_{eq} = \ell\, T^U_{eq} + b_1 \left(T^U_{eq} - T^L_{eq} \right) \quad F^L_{eq} = -b_2 \left(T^U_{eq} - T^L_{eq} \right).$$

We obtain:

$$F^U_{eq} = 312.7390 \quad F^L_{eq} = -0.0689.$$

C Additional Results

C.1 Low Damages

Figure 9: Additional Outcomes with Low Damages

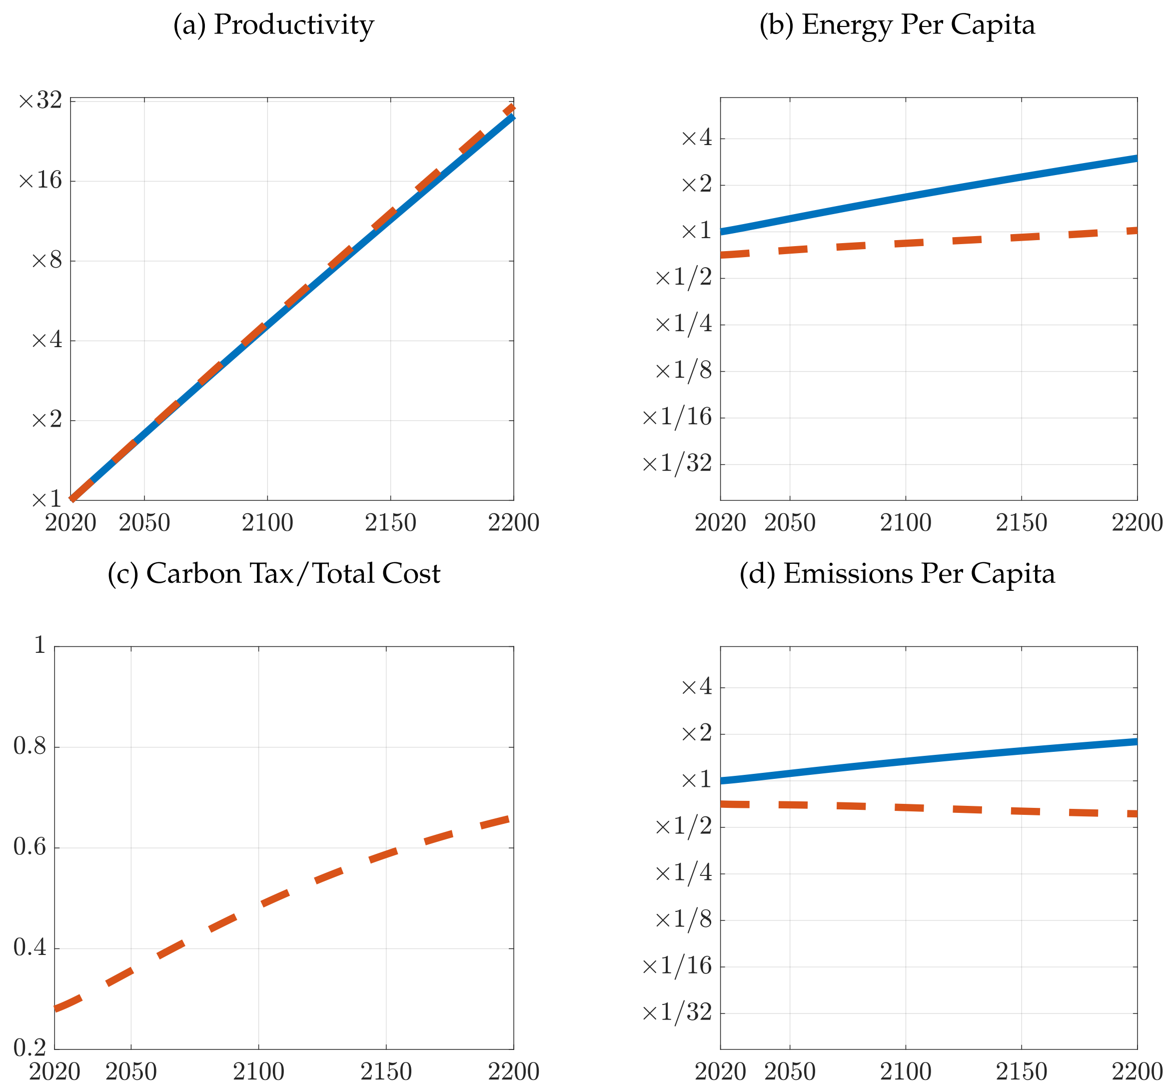

Notes: Transition paths in the laissez-faire and in the efficient allocation under the local temperature damage function ζ^{local}, estimated with local temperature variation and with cumulative impact $\int_0^\infty \zeta_s^{\text{local}} ds = 0.036$, as in Figure 4. The economy starts from 2020 initial conditions. Panel (a): log productivity. Panel (b): log energy per capita. Panel (c): optimal carbon tax as a share of the total marginal cost of energy in the efficient allocation. Panel (d): log emissions per capita. Solid blue lines: laissez-faire. Dashed red lines: efficient allocation.

C.2 High Damages

Figure 10: Additional Outcomes with High Damages

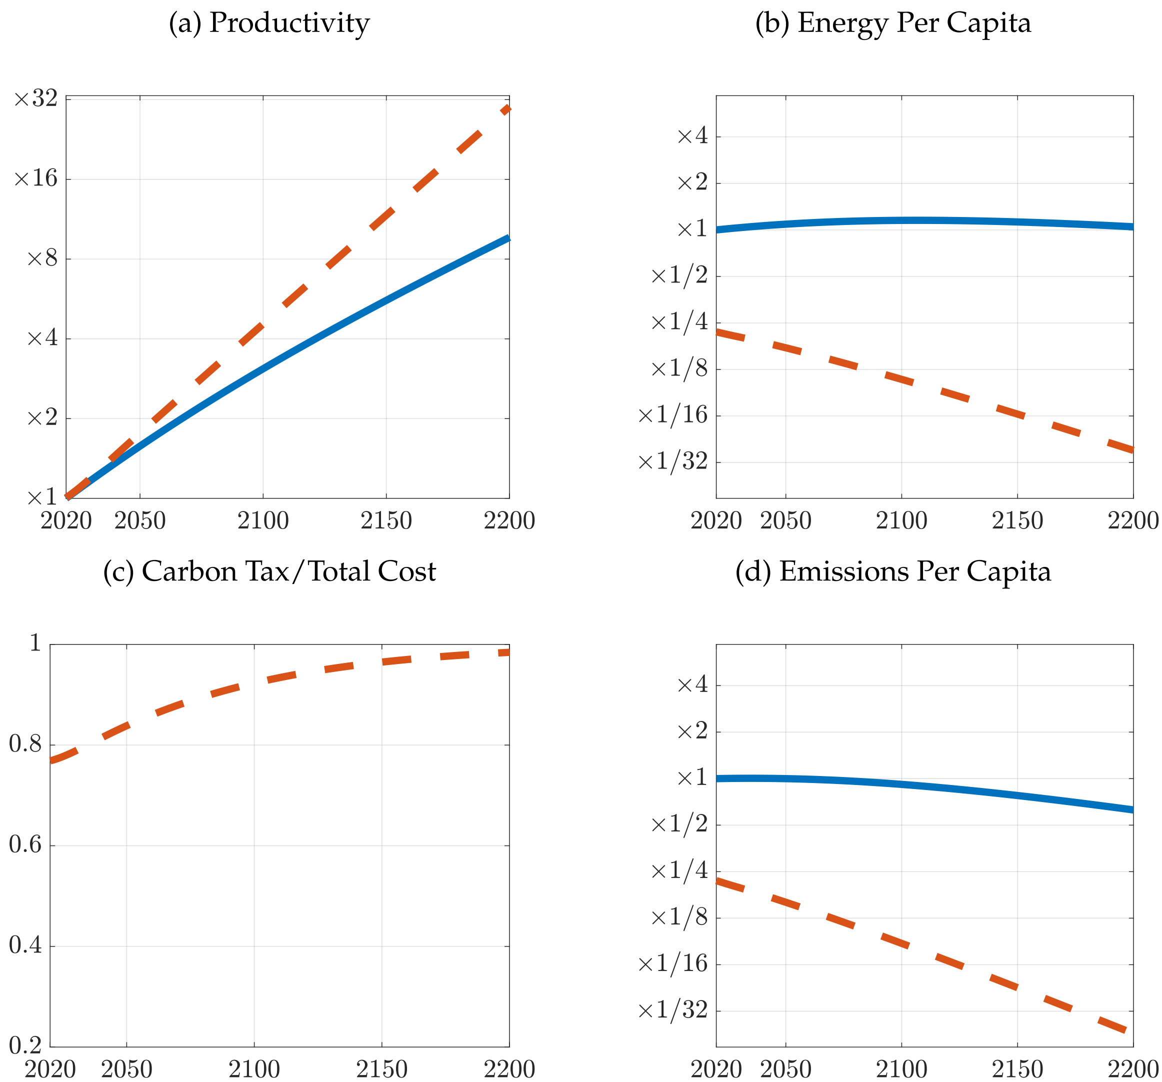

Notes: Transition paths in the laissez-faire and in the efficient allocation under the global temperature damage function ζ^{global}, estimated with global temperature variation and with cumulative impact $\int_0^\infty \zeta_s^{\text{global}} ds = 0.25$, as in Figure 5. The economy starts from 2020 initial conditions. Panel (a): log productivity. Panel (b): log energy per capita. Panel (c): optimal carbon tax as a share of the total marginal cost of energy in the efficient allocation. Panel (d): log emissions per capita. Solid blue lines: laissez-faire. Dashed red lines: efficient allocation.

C.3 Backstop Technology

Figure 11: Additional Outcomes With and Without Backstop

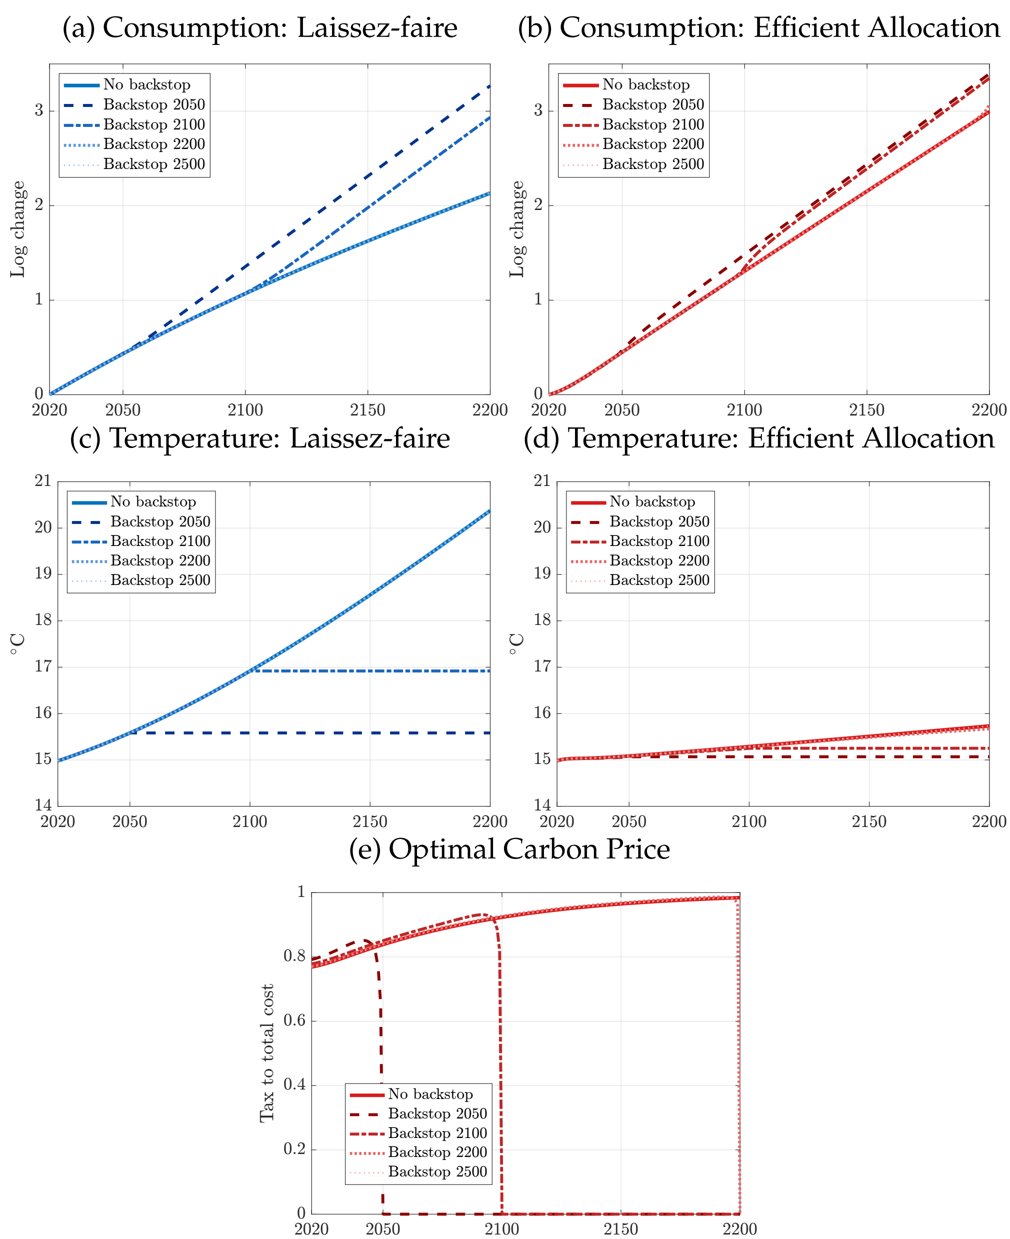

Notes: Additional outcomes when a backstop technology arrives at various dates. Damage function calibrated to global temperature variation ζ^{global}. Horizon: 2020 to 2200. Solid lines: baseline allocation without backstop. Dashed, dash-dotted, dotted and thin dotted lines: backstop arriving in 2050, 2100, 2200 and 2500. Panels (a) and (b): log change since 2020 in consumption per capita in the laissez-faire and in the efficient allocation. Panels (c) and (d): temperature in the combined atmosphere and upper ocean layer, in °C, which is constant after each arrival date. Panel (e): optimal carbon tax as a share of the total marginal cost of energy.

C.4 Nonlinear Damages

Figure 12: Marginal Damages as a Function of Warming

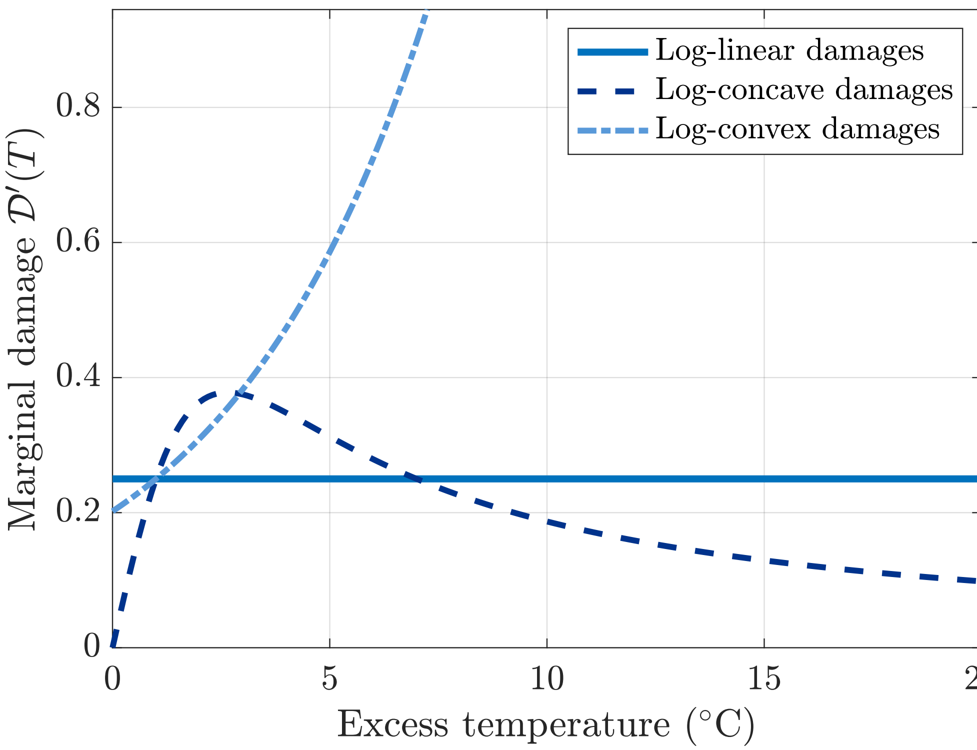

Notes: Marginal damage $D'(T)$ as a function of warming relative to 2020. Damage function calibrated to global temperature variation. All damage functions are calibrated so that marginal damages coincide at 1°C of warming above pre-industrial levels. Solid line: log-linear damage function, with constant marginal damage 0.25. Dashed line: log-concave damage function (Nordhaus, 1992; Barrage and Nordhaus, 2024). Dash-dotted line: log-convex damage function (Golosov et al., 2014).

Figure 13: Additional Outcomes With Nonlinear Damage Functions

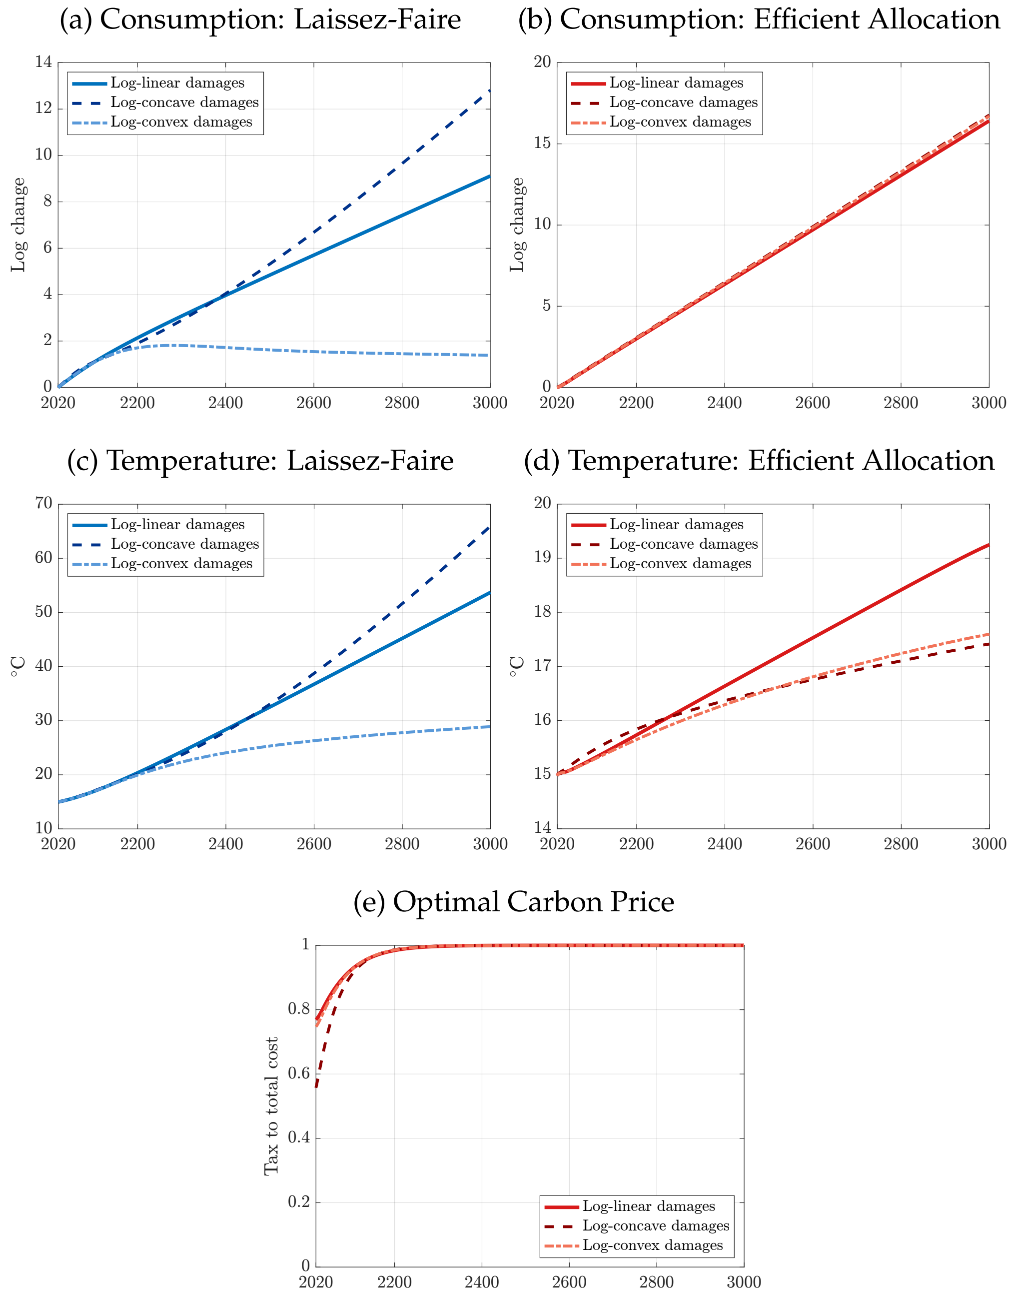

Notes: Additional outcomes under alternative damage functions, global calibration, between 2020 and 3000. Damage functions recalibrated to empirical estimates based on global temperature variation so that marginal damages coincide at 1°C of warming above pre-industrial levels. Solid lines: log-linear damages. Dashed lines: log-concave damages. Dash-dotted lines: log-convex damages. Panels (a) and (b): log change since 2020 in consumption per capita in the laissez-faire and in the efficient allocation. Panels (c) and (d): temperature in the combined atmosphere and upper ocean layer, in °C. Panel (e): optimal carbon tax as a share of the total marginal cost of energy.

Online Appendix

Balancing Growth in a Warming World

Adrien Bilal[†]
Thibault Ingrand[‡]
Diego R. Känzig[§]

Contents

[†]Stanford University, CEPR and NBER. E-mail: adrienbilal@stanford.edu.
[‡]Stanford University. E-mail: ingrand@stanford.edu.
[§]Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu.

D Data

This appendix describes the data sources that we use in the quantification of Section 3. Our estimation sample runs from 1982 to 2019. We start in 1982 to capture recent decades and exclude the oil shocks. We stop in 2019 to exclude the pandemic. The initialization of the climate system uses longer historical coverage.

Output and population. We obtain world Gross Domestic Product (GDP) and world population at an annual frequency from the World Bank World Development Indicators. World GDP is measured in constant dollars. We construct world GDP per capita as the ratio of world GDP to world population. The path of world output between 1982 and 2019 is the target of the simulated method of moments that identifies technological growth g, as described in Section 3. Average population growth over the same period identifies $n = 0.015$.

Primary energy. We obtain world primary energy consumption at an annual frequency from the Energy Institute Statistical Review of World Energy, which we retrieve through Our World in Data. The series aggregates all primary energy sources. In the model, the growth rate of the extraction cost χ_t equals the growth rate of output per unit of energy Y_t/E_t. This relationship also holds away from a balanced growth path. The average growth rate of Y_t/E_t between 1982 and 2019 therefore identifies $g_\chi = 0.012$.

Carbon emissions. We obtain carbon dioxide emissions from energy use at an annual frequency from the Potsdam Real-Time Integrated Model for the probabilistic Assessment of emission Paths (PRIMAP). The series aggregates emissions across all countries. In the model, the growth rate of emissions relative to energy use E_t^m/E_t equals the growth rate of carbon intensity σ_t, again also away from a balanced growth path. The average growth rate of E_t^m/E_t between 1982 and 2019 therefore identifies $g_\sigma = 0.003$. The long historical coverage of the emissions data also serves to initialize the climate system. We feed the observed path of emissions between 1850 and 1982 into the parametrized climate system of Appendices B and B, starting from its pre-industrial steady state. This procedure delivers the carbon stocks in all sinks and the temperatures in all layers in 1982 that serve as initial conditions in the estimation, as shown in Figure 2.

Global temperature. We obtain global mean temperature anomaly data at an annual frequency from the National Oceanic and Atmospheric Administration (NOAA, NOAA

National Centers for Environmental Information, 2023a). Temperature anomalies are deviations from the climatology, which is measured as the 1901-2000 mean temperature of 13.9°C (NOAA National Centers for Environmental Information, 2023b). Temperature levels are constructed by adding the climatology to the anomaly series. We use this data to validate the initialization of the climate system between 1850 and 1982 in Figure 2, and the fit of the estimated model between 1982 and 2019 in Figure 3.

Damage functions. The damage functions ζ_s^{local} and ζ_s^{global} are estimated in Bilal and Känzig (2026) from annual variation in country-level and global mean temperature, respectively. We refer to that paper for a detailed description of the underlying economic and climate data.

E Additional Proofs

E.1 Proof of Lemma 1

Denote $y(x) = [x\mathbf{I} - \mathbf{A}]^{-1}\mathbf{m}$. We obtain: $f(x) = \sum_i \xi_i \frac{y_i(\rho+x)}{y_i(x)}$. We characterize the limits of f by characterizing the limits of y_i.

Going back to present valuation formulas, we express: $y_i(x) = \int_0^\infty e^{-xt} v_{it} dt$, where $v_t = e^{t\mathbf{A}}\mathbf{m} > 0$. All the v_{it} are strictly positive for any $t > 0$ because $\mathbf{A}$ is irreducible. Thus, y_i is continuous.

We now show that $\log y_i(x)$ is convex. Observe that: $y_i'(x) = -\int_0^\infty te^{-tx} v_{it} dt$ and $y_i''(x) = \int_0^\infty t^2 e^{-tx} v_{it} dt$. Then: $y_i(x) y_i''(x) - (y_i'(x))^2 = \iint_0^\infty e^{-x(t+s)} v_{it} v_{is} (s^2 - ts) dt ds = \frac{1}{2} \iint_0^\infty e^{-x(t+s)} v_{it} v_{is} (t-s)^2 dt ds > 0$, where the second equality changed variables to symmetrize the expression. Thus, $\partial \log y_i(x)/\partial x$ is increasing in x.

We now use this result to characterize the monotonicity of the ratio of y_i. We obtain: $\frac{\partial}{\partial x} \log \frac{y_i(\rho+x)}{y_i(x)} = (\log y_i)'(\rho + x) - (\log y_i)'(x) > 0$. So $\log \frac{y_i(\rho+x)}{y_i(x)}$ is increasing in x, and therefore so is $\frac{y_i(\rho+x)}{y_i(x)}$. Thus, we have proven strict monotonicity of $\frac{y_i(\rho+x)}{y_i(x)}$, and, as a result, of f.

We are now ready to characterize the limits of f. $y_i(0) = +\infty$ because v_t converges to the stationary distribution of $\mathbf{A}$ in the long run. Therefore, $y_i(\rho + x)/y_i(x) \to 0$ as $x \to 0$. Thus, f goes to zero as x converges to zero. In addition, $y_i(\rho + x)/y_i(x) \to 1$ as $x \to \infty$. Thus, f converges to $\sum_i \xi_i$ as $x \to \infty$, which concludes the proof of the limits of f.

E.2 Energy Costs in Labor Units

We now suppose that the extraction cost is paid in labor units: $\chi_t E_t = L_t^E$, with $L_t = L_t^Y + L_t^E$. The only modification to our results is that emissions growth is given by $g_m = -(g_\chi + g_\sigma)$ in the laissez-faire. We prove the following result.

Corollary 1 (Growth with Labor-Intensive Energy Production)**.** *Propositions 1 and 2 apply when energy production only requires labor, with the following modifications: the regime distinction is solely based on $g_m = \min\{-(g_\chi + g_\sigma), \overline{g}_m\}$ instead of ω, and $g_Y = \frac{g - \pi g_\chi - \Delta \max\{g_m, 0\}}{1+\pi}$.*

The proof is a straightforward extension of the proof of Propositions 1 and 2.

E.2.1 Laissez-faire

An asymptotic growth path requires that the marginal product of capital is stable over time. The marginal product of capital depends on $L_t^Y = 1 - \chi_t E_t$. Therefore, $\chi_t E_t$ must stay constant or shrink over time, implying: $g_\chi + g_E \leq 0$.

Energy demand implies: $F_E = w_t \chi_t$. The household budget constraint requires that w_t grows at g_Y. Therefore: $g_E = g_Y - g_Y - g_\chi = -g_\chi$, so that $\chi_t E_t$ is exactly constant. We next obtain: $g_m = -g_\chi - g_\sigma$, which now only depends on parameters. The same distinction of cases—as when energy uses final goods—arises, where we obtain instead: $g_Y = \frac{g - \pi g_\chi - \Delta \max\{-g_\chi - g_\sigma, 0\}}{1+\pi}$.

E.2.2 Efficient Allocation

Labor optimality implies: $(1-\alpha)(1-\psi)Y_t / L_t^Y = \overline{w}_t$. Hence: $\overline{w}_t$ must grow at g_Y.

Energy demand implies: $F_E = \overline{w}_t \chi_t + \sigma_t \tau_t$, where $\overline{w}_t$ is the planner's shadow wage. Re-arranging as above: $(\psi - \phi E_t^m / M_t) \frac{Y_t}{\overline{w}_t \chi_t E_t} = 1$.

If $\psi > \phi E_t^m / M_t$, $\chi_t E_t$ must be constant. The laissez-faire formulas apply, and we need $g_m = -g_\chi - g_\sigma < \overline{g}_m$, where $\overline{g}_m$ is defined as above. If $\psi = \phi E_t^m / M_t$, $g_m = \overline{g}_m$ as above.

E.3 Exhaustible Resource Management with Exploration

We first sketch the main elements of this extension. Energy reserves R_t now evolve according to the difference between exploration X_t and energy use E_t: $\dot{R}_t = X_t - E_t$. Exploration comes at a cost that reflects the scarcity of energy: $\frac{\chi(S_t)}{2} \left(X_t / R_t\right)^2 R_t$, where S_t represents cumulative extraction and satisfies $\dot{S}_t = E_t$, and where $\chi(S_t) = \overline{\chi} S_t^{\varepsilon}$.

To focus on the carbon externality, we require that households internalize that cumulative extraction affects exploration costs. A simple microfoundation is that each household owns energy fields and understands that more extraction implies higher costs. The rest of the economy is unchanged.

To avoid a lengthy discussion of cases, we require one mild restriction that is consistent with the data: there is some progress in carbon intensity over time, $g_\sigma \geq 0$. We let $\widetilde{\varepsilon} = \frac{\varepsilon}{1+\varepsilon}$, equivalently $\varepsilon = \frac{\widetilde{\varepsilon}}{1-\widetilde{\varepsilon}}$. We prove the following result.

Corollary 2 (Growth with Energy Stock Management). *The solution of the exhaustible resource model (18) admits a unique asymptotic growth path, which is also a balanced growth path. Propositions 1, 2 and 4 apply, with the following modifications: $g_\chi = 0$, $\omega = \frac{1+\pi\widetilde{\varepsilon}}{1-\widetilde{\varepsilon}} g_\sigma$, $g_Y = \frac{g+\Delta^{ex} g_\sigma}{1+\pi\widetilde{\varepsilon}+\Delta^{ex}(1-\widetilde{\varepsilon})}$ when growth rates coincide, where $\Delta^{ex} = \Delta \mathbb{1}\{g > \omega\}$.*

E.3.1 Setup

Let R_t denote reserves of energy, which evolve according to $\dot{R}_t = X_t - E_t$, where X_t denotes exploration. Exploration comes at a cost that reflects the scarcity of energy: $\frac{\chi(S_t)}{2}\left(\frac{X_t}{R_t}\right)^2 R_t$, where S_t represents cumulative extraction and satisfies $\dot{S}_t = E_t$. We consider an isoelastic specification: $\chi(S_t) = \overline{\chi} S_t^{\frac{\widetilde{\varepsilon}}{1-\widetilde{\varepsilon}}}$, with $\frac{\widetilde{\varepsilon}}{1-\widetilde{\varepsilon}} \geq 0$. As before, p_t denotes the price of energy.

In the laissez-faire, households now solve:

$$\max_{\{C_t, K_t, E_t\}_t} \int_0^\infty e^{-\rho t} U(C_t) dt \quad \text{subject to} \quad C_t + \dot{K}_t + \frac{\chi(S_t)}{2}\left(\frac{X_t}{R_t}\right)^2 R_t = r_t K_t + w_t + p_t E_t$$
$$\text{and to} \qquad \dot{R}_t = X_t - E_t, \;\; \dot{S}_t = E_t \tag{18}$$

given K_0, R_0, S_0. To focus on the carbon externality, we require that households internalize that cumulative extraction affects exploration costs.

The rest of the economy is unchanged. We omit the definition of the efficient allocation for brevity. The shifter of the exploration cost $\overline{\chi}$ is constant because we model the effect of scarcity on extraction costs explicitly through the dependence on cumulative extraction S_t.

To avoid a lengthy discussion of cases, we require one mild restriction that is consistent with the data: there is some progress in carbon intensity over time, $g_\sigma \geq 0$. Output growth is then positive along the growth path, because $\widetilde{\varepsilon} = \frac{\varepsilon}{1+\varepsilon} \in [0,1)$ implies $\widetilde{\varepsilon} < 1$ and $1 + \pi\widetilde{\varepsilon} > 0$.

E.3.2 Exhaustible Resources with a Fixed Limit

We start with a formulation in which exploration is not possible: $X_t \equiv 0$. The partial current-value Hamiltonian (excluding the climate system) is: $\mathcal{H} = U(C_t) + \lambda_t(r_t K_t + w_t + p_t E_t - C_t) + q_t(-E_t)$. Optimality requires: $\lambda_t = u'(C_t)$, $\gamma\dot{C}_t/C_t = r_t - \rho$, $u'(C_t)p_t = q_t$, and $\rho q_t - \dot{q}_t = 0$. Combining the last two equations, we obtain: $\gamma\dot{C}_t/C_t = \dot{p}_t/p_t - \rho$. Hence we obtain the classic Hotelling result that energy prices grow at the interest rate: $\dot{p}_t/p_t = r_t$. On a constant growth path: $r = \rho + \gamma g_Y$, so that energy prices grow at rate $g_p = \rho + \gamma g_Y$.

Using that $g_m = g_Y - g_p - g_\sigma$, we obtain: $g_m = -(\rho + (\gamma-1)g_Y) - g_\sigma$, with $\rho + (\gamma-1)g_Y > 0$ by the capital transversality condition. We immediately see that $g_m < 0$ if $g_\sigma > 0$: carbon and temperature stabilize in the long run. But if $g_\sigma < 0$, we can still obtain positive emissions growth. Recall that $g_E = g_m + g_\sigma = -(\rho + (\gamma-1)g_Y) < 0$.

Case 1. $g_m < 0$. Carbon and temperature stabilize, productivity grows at rate g. Then: $(1+\pi)g_Y = g - \pi[\rho + (\gamma-1)g_Y]$. Re-arranging: $g_Y = \frac{g-\pi\rho}{1+\pi\gamma}$. We then obtain: $g_m + g_\sigma = -\frac{(1+\pi)\rho+(\gamma-1)g}{1+\pi\gamma}$. So $g_m < 0$ if and only if $g_\sigma + \frac{(1+\pi)\rho+(\gamma-1)g}{1+\pi\gamma} > 0$. This is always satisfied if $g_\sigma > 0$ and $\gamma \geq 1$.

Case 2. $g_m > 0$. Carbon grows at g_m with a productivity drag at rate Δg_m. Therefore: $(1+\pi)g_Y = g - \Delta(g_E - g_\sigma) + \pi g_E = g + (\pi - \Delta)g_E + \Delta g_\sigma = g - (\pi-\Delta)[\rho + (\gamma - 1)g_Y] + \Delta g_\sigma$. Re-arranging: $g_Y = \frac{g-\pi\rho+\Delta(\rho+g_\sigma)}{1+\pi\gamma-(\gamma-1)\Delta}$, and $g_m > 0$ based on the other side of the condition from Case 1.

E.3.3 Laissez-Faire with Exploration

We now allow for exploration, but exploration costs do not depend on cumulative past extraction: $\widetilde{\varepsilon} = 0$.

The partial current-value Hamiltonian (excluding the climate system) is: $\mathcal{H} = U(C_t) + \lambda_t(r_t K_t + w_t + p_t E_t - C_t - \chi_t(X_t/R_t)^2 R_t/2) + q_t(X_t - E_t)$. As above, optimality requires: $\lambda_t = u'(C_t)$, $\gamma\dot{C}_t/C_t = r_t - \rho$, $u'(C_t)p_t = q_t$.

The optimality condition for exploration is, after re-expressing q_t: $X_t = \frac{p_t R_t}{\chi_t}$. The co-state equation for R_t now changes. We obtain: $\rho q_t - \dot{q}_t = u'(C_t)\frac{\chi_t}{2}\left(\frac{X_t}{R_t}\right)^2$. Substituting in the other conditions: $\rho q_t - \dot{q}_t = u'(C_t)\frac{p_t^2}{2\chi_t}$. Further re-arranging: $\gamma \dot{C}_t / C_t = \frac{\dot{p}_t}{p_t} + \frac{p_t}{2\chi_t} - \rho$.

Therefore, we now obtain: $g_p + \frac{p_t}{2\chi_t} = r = \rho + \gamma g_Y$ on a constant growth path. For this condition to hold as t grows, we need the term p_t/χ_t to remain bounded. This requires $g_p \leq g_\chi$.

Case 1. $g_p = g_\chi$. We obtain exactly the same equilibrium conditions as in our baseline with a static extraction cost, with one exception. The exception is that the level of the energy price is determined by the condition $g_\chi + \frac{p}{2\chi} = \rho + \gamma g_Y$. This condition only affects the level of the energy price, not its growth rate. Thus, Proposition 1 applies.

We need to make sure that exploration costs do not affect the asymptotic constant growth path. When $g_p = g_\chi$, exploration costs grow as $\chi_t R_t$. In addition, we have: $X_t/R_t = p_t/\chi_t = 2(r - g_\chi)$. So R_t evolves according to: $\dot{R}_t = 2(r - g_\chi)R_t - E_t$. Since $g_E = g_Y - g_\chi$, we can solve explicitly: $R_t = \overline{R}e^{2(r-g_\chi)t} + \frac{E_0 e^{g_E t}}{2(r-g_\chi)-g_E}$ for some $\overline{R}$. The reserve transversality condition is $e^{-\rho t}u'(C_t)p_t R_t \to 0$, that is, $e^{-(r-g_\chi)t}R_t \to 0$. Therefore: $\overline{R} = 0$, and reserves R_t grow at rate g_E, which is consistent with the resource constraint.

Finally, we must check $g_Y < r = \rho + \gamma g_Y$. We have: $g_Y = \frac{g+(\Delta-\pi)g_\chi+\Delta g_\sigma}{1+\Delta}$. A sufficient condition for $g_Y < r$ is $(\gamma - 1)g_Y \geq 0$, i.e. $(\gamma-1)(g + (\Delta-\pi)g_\chi + \Delta g_\sigma) \geq 0$.

Case 2. $g_p < g_\chi$. We obtain $g_p = \rho + \gamma g_Y$ as in the exhaustible resource model with a hard limit. We need to check whether, indeed, $\rho + \gamma g_Y < g_\chi$. To simplify the following discussion, we impose $g_\chi = 0$ because we are already modeling scarcity directly. Hence, we would need negative output growth for this case to arise.

Case 2.1. $g_m < 0$. Then $\rho + \gamma g_Y = \frac{\rho+\gamma g}{1+\gamma\pi} > 0$, a contradiction.

Case 2.2. $g_m > 0$. Then: $g_E = g_m + g_\sigma > 0$ if $g_\sigma \geq 0$. With $X_t/R_t \to 0$ and E_t growing, reserves are depleted in finite time, a contradiction.

E.3.4 Efficient Allocation with Exploration

We still allow for exploration, but exploration costs do not depend on cumulative past extraction: $\widetilde{\varepsilon} = 0$.

The partial current-value Hamiltonian (excluding the climate system) is: $\mathcal{H} = U(C_t) + \lambda_t(F(Z(\mathbf{T}_t), K_t, 1, E_t) - \delta K_t - C_t - \chi_t(X_t/R_t)^2 R_t/2) + q_t(X_t - E_t) - \boldsymbol{\mu}_t'(\mathbf{AM}_t + \sigma_t E_t \mathbf{m}) -$

$\boldsymbol{\nu}_t'(\mathbf{BT}_t + \mathbf{NF}_{eq} + \mathbf{NH}(\log \mathbf{M}_t - \log \mathbf{M}_{eq}))$, where we normalize $\boldsymbol{\mu}_t, \boldsymbol{\nu}_t$ so that they are positive, as in Appendix A.5. Let $p_t = q_t/u'(C_t)$. The only optimality condition that differs from the laissez-faire is the one for energy: $\psi\frac{Y_t}{E_t} = p_t + \sigma_t\tau_t$, with $\tau_t = \boldsymbol{\mu}_t'\mathbf{m}/u'(C_t) \geq 0$. As in the laissez-faire, we obtain: $g_p + \frac{p_t}{2\chi_t} = \rho + \gamma g_Y$.

We consider the same cases as in the laissez-faire. If $g_p = g_\chi$, Propositions 2 and 4 apply.

If $g_p < g_\chi = 0$, we obtain the same optimality conditions as in the planner problem. Hence, Propositions 2 and 4 apply conditional on replacing g_χ by the (endogenous) expression for g_p. We also obtain $g_p = \rho + \gamma g_Y$ as in the laissez-faire. We must now verify that these Propositions apply unconditionally to g_p.

If $g_m < 0$, growth rates from the laissez-faire apply, and we obtain the same contradiction as in the laissez-faire. If $0 < g_m < \overline{g}_m$, growth rates from the laissez-faire apply and we obtain the same contradiction as in the laissez-faire.

If $g_m = \overline{g}_m > 0$, which is independent of g_p, the growth rate of p is irrelevant to the entire asymptotic growth path. Then, Propositions 2 and 4 apply directly.

E.3.5 Laissez-Faire with Exploration and Scarcity

We denote by $\varepsilon = \frac{\widetilde{\varepsilon}}{1-\widetilde{\varepsilon}}$ the elasticity of the cost shifter. The partial current-value Hamiltonian (excluding the climate system) is:

$$\mathcal{H} = U(C_t) + \lambda_t\left(r_tK_t + w_t + p_tE_t - C_t - \frac{\chi(S_t)}{2}\left(\frac{X_t}{R_t}\right)^2 R_t\right) + q_t(X_t - E_t) + v_tE_t.$$

Define $v_t = z_tu'(C_t)$ and $\widetilde{p}_t = p_t + z_t$. As above, optimality requires: $\lambda_t = u'(C_t)$, $\gamma\dot{C}_t/C_t = r_t - \rho$. Now, however, we have the energy and exploration FOCs:

$$q_t = \lambda_t\widetilde{p}_t \qquad \lambda_t\chi_t\frac{X_t}{R_t} = q_t \qquad \Rightarrow \qquad \frac{X_t}{R_t} = \frac{\widetilde{p}_t}{\chi_t}.$$

The co-state equations are as follows:

$$(R)\ \ \rho q_t - \dot{q}_t = \lambda_t\frac{\widetilde{p}_t^2}{2\chi_t} \qquad\qquad (S)\ \ \rho v_t - \dot{v}_t = -\lambda_t\frac{\chi_t'}{2\chi_t^2}\widetilde{p}_t^2R_t$$

where $\chi'_t = \chi'(S_t)$. Using the Euler equation: $\dot{q} = \lambda((\rho - r)\widetilde{p} + \dot{\widetilde{p}})$, and $\dot{v} = \lambda((\rho - r)z + \dot{z})$. Therefore:

$$r_t\widetilde{p}_t - \dot{\widetilde{p}}_t = \frac{\widetilde{p}_t^2}{2\chi_t} \qquad\qquad r_t z_t - \dot{z}_t = -\frac{\chi'_t}{2\chi_t^2}\widetilde{p}_t^2 R_t$$

On a constant growth path, the first equation implies:

$$r = g_{\widetilde{p}} + \frac{\widetilde{p}_t}{2\chi_t}.$$

Now let $\chi(S) = \overline{\chi}S^\varepsilon$ have a constant elasticity. Then: $g_\chi = \varepsilon g_S = \varepsilon g_E = \varepsilon(g_Y - g_p)$ where the second to last equality follows from the fact that S_t grows at the growth rate of E_t, and the last equality follows from the energy FOC of the firm. The z_t equation becomes:

$$rz_t - \dot{z}_t = -\frac{\varepsilon}{2}\frac{\widetilde{p}_t^2 R_t}{\chi_t S_t}.$$

We now distinguish cases as before.

Case 1. $g_{\widetilde{p}} = g_\chi$. Then: $\widetilde{p}/2\chi$ is a constant. We obtain: $g_{\widetilde{p}} = g_\chi = \varepsilon g_E$. The z_t equation then becomes: $rz_t - \dot{z}_t = -\xi\frac{\widetilde{p}_t R_t}{S_t}$, with $\xi = \frac{\varepsilon}{2}\frac{\widetilde{p}}{\chi}$.

The term $\frac{\widetilde{p}_t R_t}{S_t}$ grows at rate $(\varepsilon - 1)g_E + g_R$. Solving the z_t equation forward, we obtain that z_t is negative and grows at rate $g_z = (\varepsilon - 1)g_E + g_R$, with $g_E = g_Y - g_p$.

As before, we obtain that $\dot{R}_t = 2(r - \varepsilon g_E)R_t - E_t$. So R_t grows at $g_R = \max\{2(r - \varepsilon g_E), g_E\}$. The R transversality condition implies that $e^{-\rho t}q_t R_t \to 0$. Re-arranging, this implies: $e^{-rt}\widetilde{p}_t R_t \to 0$. Since $\widetilde{p}_t$ grows at εg_E, the TVC term grows at $\max\{r - \varepsilon g_E, (1 + \varepsilon)g_E - r\}$. Therefore, the component of R_t that grows at $2(r - \varepsilon g_E)$ must be zero, and so: $g_R = g_E$. We will need to verify whether, in equilibrium, $(1 + \varepsilon)g_E < r$ and $\varepsilon g_E < r$.

Since $g_R = g_E$, we obtain that $g_z = \varepsilon g_E = g_\chi$. Thus, $\widetilde{p}$ and z grow at the same rate, so that $p_t = \widetilde{p}_t - z_t$ with $\widetilde{p}_t$ and $-z_t$ both positive implies: $g_p = g_{\widetilde{p}} = g_z = g_\chi = \varepsilon g_E$. The firm's energy FOC $\psi Y_t/E_t = p_t$ then implies: $g_Y = g_E + g_p = (1 + \varepsilon)g_E$, and so $g_\chi = \frac{\varepsilon}{1+\varepsilon}g_Y$. Note that $(1 + \varepsilon)g_E = g_Y < r$ holds by the capital transversality condition, and $\varepsilon g_E = g_\chi < g_Y < r$ given $g_Y > 0$, which we verify in the cases below. This confirms the requirements above.

We can now apply Proposition 1 conditional on $g_\chi = \frac{\varepsilon}{1+\varepsilon}g_Y$. To simplify the discussion, we define: $\widetilde{\varepsilon} = \frac{\varepsilon}{1+\varepsilon}$, which spans $[0, 1)$ for $\varepsilon \geq 0$. We have: $g_\chi = \widetilde{\varepsilon}g_Y$.

Case 1.1. $g < \omega$. Then $g_m < 0$, and case (i) applies: $g_Y = g - \pi g_\chi$, implying $g_Y = \frac{g}{1+\pi\widetilde{\varepsilon}}$. The condition for $g_m < 0$ becomes: $g_Y < g_\chi + g_\sigma$, i.e.: $(1-\widetilde{\varepsilon})g_Y < g_\sigma$.

Since $\widetilde{\varepsilon} \in [0,1)$, we have $1 + \pi\widetilde{\varepsilon} > 0$, so that $g_Y > 0$. Then, the condition for $g_m < 0$ requires $\frac{1-\widetilde{\varepsilon}}{1+\pi\widetilde{\varepsilon}} g < g_\sigma$, where both factors on the left-hand side are positive.

Case 1.2. $g > \omega$. Then $g_m > 0$, and case (iii) applies: $g_Y = g_\chi + g_\sigma + \frac{g-\omega}{1+\Delta} = \frac{g+(\Delta-\pi)g_\chi+\Delta g_\sigma}{1+\Delta}$, implying $g_Y = \frac{g+\Delta g_\sigma}{1+\Delta+\widetilde{\varepsilon}(\pi-\Delta)}$. The condition for $g_m > 0$ is the converse of the one from Case 1.1.

Since $\widetilde{\varepsilon} \in [0,1)$, the denominator satisfies $1 + \Delta + \widetilde{\varepsilon}(\pi - \Delta) = 1 + \pi\widetilde{\varepsilon} + \Delta(1-\widetilde{\varepsilon}) > 0$, so that $g_Y > 0$ given $g_\sigma \geq 0$.

Case 2. $g_{\widetilde{p}} < g_\chi$. We then obtain $g_{\widetilde{p}} = r = \rho + \gamma g_Y$. As above, we use the resource accumulation equation to characterize g_p. $\dot{R}_t/R_t = X_t/R_t - E_t/R_t$. Since $X_t/R_t \to 0$ by assumption in Case 2, $g_R = g_E < 0$.

Since $g_E < 0$, cumulative extraction S_t converges, so that $g_\chi = \varepsilon g_S = 0$ and this case requires $r = g_{\widetilde{p}} < g_\chi = 0$. Temperature stabilizes because $g_m = g_E - g_\sigma < 0$, so $g_Z = g$, and the production function together with the energy FOC gives $g_Y = g - \pi r > 0$. Then $r = \rho + \gamma g_Y > 0$, a contradiction.

E.3.6 Efficient Allocation with Exploration and Scarcity

The only optimality condition that differs from the laissez-faire is energy use: $\psi\frac{Y_t}{E_t} = \widetilde{p}_t - z_t + \sigma_t\tau_t$, with $\tau_t = \boldsymbol{\mu}_t'\mathbf{m}/u'(C_t) \geq 0$: the laissez-faire condition $p_t = \widetilde{p}_t - z_t$ augmented by the carbon tax.

If $g_{\widetilde{p}} = g_\chi$, Propositions 2 and 4 apply with the same modifications as for the laissez-faire since $\overline{g}_m$ is determined by an equation where the only dependence on the economic module is g_Y.

If $g_{\widetilde{p}} < g_\chi$: if $g_m < \overline{g}_m$, growth rates from the laissez-faire apply, and we obtain the same contradiction. If $g_m = \overline{g}_m$, which is independent of $g_{\widetilde{p}}$, the growth path is independent of $g_{\widetilde{p}}$, and Propositions 2 and 4 apply.

E.4 Endogenous Technological Growth

We first sketch the main features of this extension, with a complete description below. We enrich our framework with endogenous growth. Instead of exogenous technological growth g, in this section we microfound the creative destruction process through innovation as in Grossman and Helpman (1991) and Aghion and Howitt (1992).

A Cobb-Douglas aggregator of varieties now constitutes aggregate output. Each variety is produced with the production function (1) subject to a variety-specific productivity. Within varieties, there is perfect competition. At any given point in time for any variety, a leader enjoys a productivity advantage $\lambda > 1$ over its competitors, leading to limit pricing. Firms use labor to innovate over the leader, with one unit of labor leading to a success rate of η. Labor is perfectly mobile between production and innovation.

The only difference from our baseline specification is that the allocation of scarce labor between innovation and production determines the growth rate of technology g. The relationship $g_Z = g - \Delta \max\{g_m, 0\}$ continues to hold, where g is a function of parameters: the rate of time preference ρ, the innovation step size λ, the innovation success rate per unit of labor η and the labor elasticity of production.

We formalize this discussion by proving the following result.

Corollary 3 (Endogenous Technological Growth)**.** *Propositions 1, 2 and 4 apply to the endogenous growth model, with the following modifications. For $i = LF, SP$, the threshold for growing temperatures is given by an i-specific carbon savings growth rate $\omega^i = (1 + \Gamma^i + \pi)g_\chi + (1 + \Gamma^i)g_\sigma$ and an i-specific technological growth rate g^i, where $\Delta^i = \Delta \mathbb{1}\{g^i > \omega^i\}$ and Γ^i is proportional to $\gamma - 1$. In addition, the growth rate of output is $g_Y^{LF} = \frac{g^{LF} + (\Delta^{LF} - \pi)g_\chi + \Delta^{LF} g_\sigma}{1 + \Gamma^{LF} + \Delta^{LF}}$ in the laissez-faire. In the efficient allocation: $g_Y^{SP} = \frac{g^{SP} + (\Delta^{SP} - \pi)g_\chi + \Delta^{SP} g_\sigma}{1 + \Gamma^{SP} + \Delta^{SP}}$ if $g^{SP} < \omega^{SP} + (1 + \Gamma^{SP} + \Delta)\overline{g}_m$, and $g_Y^{SP} = \frac{g^{SP} + (\pi - \Delta)\overline{g}_m + \pi g_\sigma}{1 + \Gamma^{SP} + \pi}$ otherwise.*

E.4.1 Setup

Consider a quality ladder model as in Grossman and Helpman (1991) and Aghion and Howitt (1992). There is now a continuum of intermediate varieties $j \in [0, 1]$, across which households have Cobb-Douglas preferences with uniform weights. They all use the production function F, but may differ in productivity/quality. We have:

$$\log Y_t = \int_0^1 \log y_{jt} dj.$$

Demand for each variety is given by: $p_{jt} y_{jt} = Y_t$, where we use the aggregate price index as the numeraire. For a monopolist in variety j, the unconstrained profit-maximizing price would therefore entail an infinite markup over marginal cost.

Each variety j is produced according to:

$$y_{jt} = z_{jt}^{(1-\alpha)(1-\psi)} F_t^* \ell_{jt},$$

where F_t^* denotes the production function evaluated at aggregate productivity of 1 and optimal input ratios $k_{jt}/\ell_{jt}, e_{jt}/\ell_{jt}$, ℓ_{jt} denotes labor demand, and z_{jt} is productivity of variety j.

There is free entry and perfect competition in each variety. In each variety, there is always a leader with a productivity advantage $z_{jt} = \lambda \widetilde{z}_{jt}$ over a competitive fringe with the old technology $\widetilde{z}_{jt}$. Let $\overline{z}_{jt} = \widetilde{z}_{jt}^{(1-\alpha)(1-\psi)}$. Given perfect competition, the leader uses limit pricing: $p_{jt} = \mathrm{mc}_{jt} = \frac{w_t}{\overline{z}_{jt} F_t^*}$. Quantity is then: $y_{jt} = \frac{Y_t \overline{z}_{jt} F_t^*}{w_t}$. Profits are: $\pi_{jt} = \frac{\overline{\lambda}-1}{\overline{\lambda}} Y_t \equiv \varpi Y_t$, with $\overline{\lambda} = \lambda^{(1-\alpha)(1-\psi)}$.

Innovation labor n_{jt} allows entrants to innovate on the leader by the step size λ at rate $\iota_{jt} = \eta n_{jt}$. The value of holding the patent to the leading technology satisfies:

$$r_t V_{jt} - \dot{V}_{jt} = \pi_{jt} - \iota_{jt} V_{jt},$$

where the last term represents displacement by new entrants. Free entry into innovation requires: $w_t = \eta V_{jt}$.

We focus on a symmetric equilibrium. We obtain: $(r_t + \iota_t)V_t - \dot{V}_t = \varpi Y_t$. Therefore, on a constant growth path, $g_w = g_V = g_Y$. Hence: $(r + \iota - g_Y)V_t = \varpi Y_t$, and $w_t = \eta V_t$. Re-arranging:

$$\frac{w_t}{Y_t} = \frac{\eta \varpi}{r + \iota - g_Y} = \frac{\eta \varpi}{\rho + \iota + (\gamma - 1) g_Y}.$$

Using the production FOC w.r.t. labor, and denoting $\kappa = (1-\alpha)(1-\psi)$, we obtain:

$$\kappa = \frac{\eta \varpi L^Y}{\rho + \iota + (\gamma - 1) g_Y}.$$

ι is determined by the amount of labor in innovation:

$$\iota = \eta L^I = \eta(1 - L^Y).$$

Substituting into the previous equation, we pin down L^Y from:

$$\kappa = \frac{\eta\varpi L^Y}{\rho + \eta(1 - L^Y) + (\gamma - 1)g_Y}.$$

Its solution is

$$L^Y = L_0^Y + L_1^Y(\gamma - 1)g_Y,$$

where L_0^Y and L_1^Y are functions of parameters only.

Finally, since production admits constant returns, production is isomorphic to an aggregate production function with aggregate productivity Z_t:

$$\log Z_t = \int_0^1 \log z_{jt} dj.$$

For every period dt, a fraction ιdt of varieties experience a successful innovation, so that Z_t grows at exponential rate:

$$g_Z = \iota \log \lambda = \eta(1 - L^Y) \log \lambda = g - \Gamma g_Y$$

where $g = \eta(1 - L_0^Y) \log \lambda$, and $\Gamma = (\gamma - 1)\eta L_1^Y \log \lambda$. Under log utility, technology growth is entirely determined as a function of parameters.

E.4.2 Laissez-Faire

Under log utility, Proposition 1 applies directly. Away from log utility, we only have to account for the equilibrium feedback between g_Z and g_Y. Recall that $g_m = g_Y - g_\chi - g_\sigma$.

Case 1. $\boldsymbol{g_m < 0}$. Then: $g_Y = g_Z - \pi g_\chi = g - \Gamma g_Y - \pi g_\chi$, so that: $g_Y = \frac{g - \pi g_\chi}{1+\Gamma}$. Then: $g_m = \frac{g-\pi g_\chi}{1+\Gamma} - g_\chi - g_\sigma$.

Case 2. $\boldsymbol{g_m > 0}$. Then: $g_Y = g_Z - \pi g_\chi = g - \Gamma g_Y - \Delta(g_Y - g_\chi - g_\sigma) - \pi g_\chi$, so that: $g_Y = \frac{g+(\Delta-\pi)g_\chi + \Delta g_\sigma}{1+\Gamma+\Delta}$.

E.4.3 Efficient Allocation

Planner conditions. We set up the planning problem after aggregation. The planner solves:

$$\begin{aligned} & \max_{\{C_t, K_t, E_t, L_t^Y\}_t} \int_0^\infty e^{-\rho t} U(C_t) dt \\ \text{subject to} \quad & C_t + \dot{K}_t = F(Z_t, K_t, L_t^Y, E_t) - \chi_t E_t \\ & \frac{\dot{Z}_t}{Z_t} = \eta(1 - L_t^Y) \log \lambda \\ \text{and to} \quad & \text{equations (2), (3), (6), (5), (7)} \end{aligned} \tag{19}$$

The partial current-value Hamiltonian (omitting the climate module) is:

$$\mathcal{H} = U(C_t) + \lambda_t \Big(F(Z_t, K_t, L_t^Y, E_t) - \chi_t E_t - C_t \Big) + q_t Z_t \eta (1 - L_t^Y) \log \lambda.$$

The optimality condition for production labor and the co-state equation for Z_t are:

$$F_{Lt} = \eta \frac{q_t}{\lambda_t} Z_t \log \lambda \qquad \qquad \rho q_t - \dot{q}_t = \lambda_t F_{Zt} + q_t \eta (1 - L_t^Y) \log \lambda.$$

Denote $q_t = \frac{\lambda_t V_t}{Z_t}$ and $\iota_t = \eta(1 - L_t^Y)$, so that re-arranging the first equation, and combining the second equation with the Euler equation:

$$\kappa = \eta \frac{V_t}{Y_t} L_t^Y \log \lambda \qquad \qquad r_t V_t - \dot{V}_t = \kappa Y_t - g_Z V_t + \iota_t V_t \log \lambda.$$

where $r_t = F_K$. Now, by definition, g_Z cancels out with the last term. In addition, on a constant growth path, V_t and Y_t grow at g_Y, so that:

$$(r - g_Y) V_t = \kappa Y_t.$$

Combining with the FOC, we obtain:

$$1 = \frac{\eta \log \lambda}{\rho + (\gamma - 1) g_Y} L^Y.$$

This implies:

$$g_Z = \eta \log \lambda - [\rho + (\gamma - 1) g_Y].$$

Under log utility, $g_Z = \eta \log \lambda - \rho$ is once more pinned down by parameters.

Emissions. Under log utility, Proposition 2 applies directly, albeit with a different $g = \eta \log \lambda - \rho$ than in the decentralized equilibrium. Away from log utility, we only have to account for the equilibrium feedback between g_Z and g_Y. Recall that $g_m = g_Y - g_\chi - g_\sigma$. Define $\Gamma = \gamma - 1$ in the planner problem.

Case 1. $g_m < 0$. Then: $g_Y = g_Z - \pi g_\chi = g - \Gamma g_Y - \pi g_\chi$, so that: $g_Y = \frac{g - \pi g_\chi}{1+\Gamma}$. Then: $g_m = \frac{g - \pi g_\chi}{1+\Gamma} - g_\chi - g_\sigma$.

Case 2. $\overline{g}_m > g_m > 0$. $\overline{g}_m$ satisfies the same equation as in the baseline economy, since it is independent of g_Z. Then: $g_Y = g_Z - \pi g_\chi = g - \Gamma g_Y - \Delta(g_Y - g_\chi - g_\sigma) - \pi g_\chi$, so that: $g_Y = \frac{g + (\Delta - \pi) g_\chi + \Delta g_\sigma}{1+\Gamma+\Delta}$.

Case 3. $g_m = \overline{g}_m$. Then: $(1+\pi) g_Y = g_Z + \pi(\overline{g}_m + g_\sigma) = g - \Gamma g_Y - \Delta \overline{g}_m + \pi(\overline{g}_m + g_\sigma)$, so that: $g_Y = \frac{g + (\pi - \Delta)\overline{g}_m + \pi g_\sigma}{1+\Gamma+\pi}$.

Proposition 4 applies conditional on these expressions for growth rates.

E.5 Endogenous Progress in Renewables

We first sketch the main elements of this extension, before detailing the setup below. We enrich our framework with endogenous growth in renewables. Instead of exogenous growth in χ_t and σ_t, in this section we introduce innovation in renewables in the tradition of Grossman and Helpman (1991) and Aghion and Howitt (1992), similarly to above when we introduced it for general technological growth. We use a formulation in which energy is a bundle of carbon-free and carbon-intensive sources with a unit elasticity of substitution, as in Appendix A.1.

The only difference from our baseline specification is that the allocation of scarce labor between innovation in renewables and production determines the growth rate of productivity of renewables. We denote by $\overline{g}_\chi, \overline{g}_\sigma$ the exogenous growth rates of extraction costs and emissions intensity if renewables had constant productivity, and let g_χ, g_σ denote the endogenous growth rates of extraction costs and emissions intensity.

Innovation endogenizes g_χ and g_σ through market forces. As above however, g_χ and g_σ are entirely determined by parameters that reflect the rate of time preference, the innovation step size, the innovation success rate per unit of labor, and output elasticities. These rates now differ between the laissez-faire and the efficient allocation, which we denote with superscripts LF and SP.

We consider the energy bundle from Appendix A.2 and embed endogenous growth in the renewable energy sector. Recall that we then have:

$$\chi_t = \overline{\chi}_t \xi_t^\beta \qquad \sigma_t = \overline{\sigma}_t \xi_t^\beta$$
$$g_\chi = \overline{g}_\chi + \beta g_\xi \qquad g_\sigma = \overline{g}_\sigma - \beta g_\xi,$$

where ξ_t is the extraction cost of green energy, and $\overline{\chi}_t, \overline{\sigma}_t$ represent extraction costs of fossil fuel energy. β is the Cobb-Douglas energy share of green energy. Recall also that with our convention we define $\sigma_t = \sigma_0 e^{-g_\sigma t}$.

Renewable energy is produced with the final good with relative productivity $1/\xi_t$. As above, there is a unit Cobb-Douglas continuum of varieties of renewable goods, each facing a quality ladder. We can directly apply the formulas above. We thus prove the following result.

Corollary 4 (Endogenous Progress in Renewable Energy)**.** *Propositions 1, 2 and 4 apply to the model with innovation in renewables, with the following modifications. For $i = LF, SP$, the threshold for growing temperatures is given by an i-specific carbon savings growth rate $\omega^i = \overline{g}_\chi + \pi\widetilde{g}^i_\chi + \pi\Gamma^i\overline{g}_\chi + (1+\pi\Gamma^i)\overline{g}_\sigma$, where $\Gamma^i, \widetilde{g}^i_\chi, \widetilde{g}^i_\sigma$ are defined in the proof, and Γ^i are proportional to $\gamma - 1$. Further define $\Delta^i = \Delta\mathbb{1}\{g > \omega^i\}$. In the laissez-faire, the growth rate of output is $g^{LF}_Y = \frac{g+\Delta^{LF}\overline{g}_\chi - \pi\widetilde{g}^{LF}_\chi + \Delta^{LF}\overline{g}_\sigma}{1+\pi\Gamma^{LF}+\Delta^{LF}}$. In the efficient allocation: $g^{SP}_Y = \frac{g+\Delta^{SP}\overline{g}_\chi - \pi\widetilde{g}^{SP}_\chi + \Delta^{SP}\overline{g}_\sigma}{1+\pi\Gamma^{SP}+\Delta^{SP}}$ if $g < \omega^{SP} + (1+\pi\Gamma^{SP}+\Delta)\overline{g}_m$, and $g^{SP}_Y = \frac{g+(\pi-\Delta)\overline{g}_m + \pi\widetilde{g}^{SP}_\sigma}{1+\pi\Gamma^{SP}+\pi}$ otherwise. Emissions growth follows similarly adjusted formulas given in the proof.*

E.5.1 Laissez-Faire

On a growth path, we obtain:

$$\frac{w_t}{\xi_t G_t} = \frac{\eta\varpi}{\rho + \iota + (\gamma - 1) g_Y} \text{ iff } L_t^Y \in (0,1)$$
$$\frac{w_t}{\xi_t G_t} > \frac{\eta\varpi}{\rho + \iota + (\gamma - 1) g_Y} \text{ iff } L_t^Y = 1$$
$$\frac{w_t}{\xi_t G_t} < \frac{\eta\varpi}{\rho + \iota + (\gamma - 1) g_Y} \text{ iff } L_t^Y = 0$$

where now market size is spending on green energy $\xi_t G_t$ instead of Y_t (note that when market size was Y_t, there was implicitly a unit price of the numeraire). $\xi_t G_t$ is the total number of final good units needed to produce G_t green energy. Innovation is still paid in labor units. Using the production FOCs for labor and for energy, we obtain:

$$\kappa = \frac{w_t L_t^Y}{Y_t} \qquad \psi = \frac{\chi_t E_t}{Y_t} \qquad \beta = \frac{\xi_t G_t}{\chi_t E_t}.$$

Combining these equations, we obtain:

$$\frac{w_t}{\xi_t G_t} = \frac{\kappa}{\beta\psi}\frac{1}{L_t^Y} \qquad \beta\psi = \frac{\xi_t G_t}{Y_t},$$

with the latter equation implying $g_\xi + g_G = g_Y$.

We start by assuming an interior solution. Then:

$$\frac{\kappa}{\beta\psi} = \frac{\eta\varpi L^Y}{\rho + \eta(1 - L^Y) + (\gamma - 1) g_Y}.$$

As above, the solution of this equation is:

$$L^Y = L_0^Y + L_1^Y(\gamma - 1) g_Y,$$

where L_0^Y, L_1^Y are functions of parameters only. The definition of the growth rate of productivity in the renewable sector is:

$$-g_\xi = \eta(1 - L_t^Y)\log\lambda.$$

Therefore, we obtain:

$$g_\xi = -\eta \log \lambda \Big(1 - L_0^Y - L_1^Y(\gamma-1)g_Y\Big) \equiv g_{\xi,0} + \Gamma_\xi g_Y \qquad \Gamma_\xi \equiv (\gamma-1)L_1^Y \eta \log \lambda.$$

This equation then implies:

$$\begin{aligned} g_\chi &= \widetilde{g}_\chi + \Gamma g_Y & \qquad \widetilde{g}_\chi &= \overline{g}_\chi + \beta g_{\xi,0} & \qquad \Gamma &= \beta(\gamma-1)L_1^Y \eta \log \lambda \\ g_\sigma &= \widetilde{g}_\sigma - \Gamma g_Y & \qquad \widetilde{g}_\sigma &= \overline{g}_\sigma - \beta g_{\xi,0} \end{aligned}$$

As above, with log utility, Proposition 1 applies directly with g_χ, g_σ now functions of parameters.

Away from log utility we have to incorporate the equilibrium link between g_χ, g_σ and g_Y. Recall that $g_m = g_Y - g_\chi - g_\sigma = g_Y - \widetilde{g}_\chi - \widetilde{g}_\sigma = g_Y - \overline{g}_\chi - \overline{g}_\sigma$. With a Cobb-Douglas energy bundle, technological progress in renewables exactly cancels out with the higher implied fossil fuel demand in combined carbon savings σ_t/χ_t.

Case 1. $g_m < 0$. Then: $g_Y = g_Z - \pi g_\chi = g - \pi \widetilde{g}_\chi - \pi \Gamma g_Y$, so that:

$$g_Y = \frac{g - \pi \widetilde{g}_\chi}{1 + \pi \Gamma}.$$

Then: $g_m = \frac{g - \pi \widetilde{g}_\chi}{1+\pi\Gamma} - \overline{g}_\chi - \overline{g}_\sigma$. $g_m < 0$ then is equivalent to:

$$g < \overline{g}_\chi + \pi \widetilde{g}_\chi + \pi \Gamma \overline{g}_\chi + (1 + \pi\Gamma)\overline{g}_\sigma \equiv \omega,$$

assuming $1 + \pi\Gamma > 0$.

Case 2. $g_m > 0$. Then: $g_Y = g_Z - \Delta g_m - \pi g_\chi = g - \Delta\big(g_Y - \overline{g}_\chi - \overline{g}_\sigma\big) - \pi \widetilde{g}_\chi - \pi \Gamma g_Y$, so that:

$$g_Y = \frac{g + \Delta \overline{g}_\chi - \pi \widetilde{g}_\chi + \Delta \overline{g}_\sigma}{1 + \pi\Gamma + \Delta}.$$

E.5.2 Efficient Allocation

The partial current-value Hamiltonian (omitting the climate module) is:

$$\mathcal{H} = U(C_t) + \lambda_t \Big(F(Z_t, K_t, L_t^Y, E_t) - \overline{\chi}_t \xi_t^\beta E_t - C_t\Big) - q_t \xi_t \eta (1 - L_t^Y) \log \lambda,$$

where we have already optimized the choice of green vs. brown energy. The optimality condition for production labor and the co-state equation for ξ_t are:

$$F_{Lt} = -\eta \log \lambda \frac{q_t}{\lambda_t}\xi_t \qquad\qquad \rho q_t - \dot{q}_t = -\beta\lambda_t \frac{\chi_t E_t}{\xi_t} - q_t \eta(1-L_t^Y)\log\lambda.$$

Denote $q_t = \frac{\lambda_t V_t}{\xi_t}$ and $\iota_t = \eta(1-L_t^Y)$, so that re-arranging the first equation, and combining the second equation with the Euler equation:

$$\kappa = -\eta \frac{V_t}{Y_t} L_t^Y \log\lambda \qquad\qquad r_t V_t - \dot{V}_t = -\beta\psi Y_t - g_\xi V_t - \iota_t V_t \log\lambda.$$

where $r_t = F_K$. Now, by definition, g_ξ cancels out with the last term. In addition, on a constant growth path, V_t and Y_t grow at g_Y, so that:

$$(r - g_Y)V_t = -\beta\psi Y_t.$$

Combining with the FOC, we obtain:

$$\kappa = \frac{\beta\psi\eta\log\lambda}{\rho + (\gamma-1)g_Y} L^Y.$$

Using $\kappa/(\beta\psi) = 1/(\beta\pi)$, this implies:

$$\begin{aligned}
g_\xi &= -\left\{\eta\log\lambda - \frac{\rho}{\beta\pi}\right\} + \frac{(\gamma-1)}{\beta\pi} g_Y \equiv g_{\xi,0} + g_{\xi,1} g_Y \\
g_\chi &= \overline{g}_\chi + \beta g_{\xi,0} + \frac{\gamma-1}{\pi} g_Y \equiv \widetilde{g}_\chi + \Gamma g_Y \\
g_\sigma &= \overline{g}_\sigma - \beta g_{\xi,0} - \Gamma g_Y \equiv \widetilde{g}_\sigma - \Gamma g_Y.
\end{aligned}$$

where $g_{\xi,0} = -\left\{\eta\log\lambda - \frac{\rho}{\beta\pi}\right\}$ and $\Gamma = \frac{\gamma-1}{\pi}$.

Emissions. Under log utility, Proposition 2 applies directly, albeit with different g_χ, g_σ than in the decentralized equilibrium. Away from log utility, we only have to account for the equilibrium feedback between g_χ, g_σ and g_Y. Recall that $g_m = g_Y - g_\chi - g_\sigma = g_Y - \overline{g}_\chi - \overline{g}_\sigma$.

The same formulas as in the laissez-faire apply for $g_m < \overline{g}_m$, conditional on changing Γ and $\widetilde{g}_\chi, \widetilde{g}_\sigma$ appropriately.

Case 1. $g_m < 0$. We obtain: $g_Y = \frac{g - \pi\widetilde{g}_\chi}{1+\pi\Gamma}$ and $g_m < 0$ iff $g < \overline{g}_\chi + \pi\widetilde{g}_\chi + \pi\Gamma\overline{g}_\chi + (1 + \pi\Gamma)\overline{g}_\sigma \equiv \omega$ assuming $1 + \pi\Gamma > 0$.

Case 2. $\overline{g}_m > g_m > 0$. $\overline{g}_m$ satisfies the same equation as in the baseline economy, since it is independent of g_χ, g_σ. Then the laissez-faire formula applies conditional on $\widetilde{g}_\chi, \widetilde{g}_\sigma, \Gamma$: $g_Y = \frac{g + \Delta\overline{g}_\chi - \pi\widetilde{g}_\chi + \Delta\overline{g}_\sigma}{1+\pi\Gamma+\Delta}$.

Case 3. $g_m = \overline{g}_m$. Then: $(1+\pi)g_Y = g_Z + \pi(\overline{g}_m + g_\sigma) = g - \Delta\overline{g}_m + \pi\overline{g}_m + \pi\widetilde{g}_\sigma - \pi\Gamma g_Y$, so that: $g_Y = \frac{g + (\pi - \Delta)\overline{g}_m + \pi\widetilde{g}_\sigma}{1+\Gamma\pi+\pi}$.

Proposition 4 applies conditional on these expressions for growth rates.

F Climate Module and Thermohaline Circulation

The warming module from Folini et al. (2024) is specified in deviations from pre-industrial values. Once specified in levels, it is isomorphic to our specification with exogenous forcing in the deep oceans. This exogenous forcing is necessary to ensure that, in equilibrium, the deep oceans are colder than the upper ocean and atmospheric layer as in the data.

In practice, the colder temperature of deep oceans is in part caused by the thermohaline circulation (the "global conveyor belt"). The thermohaline circulation is caused by spatial heterogeneity in atmospheric temperatures and upper ocean temperatures together with salinity. Water cools at the poles where the atmosphere is colder (and local radiative forcing is weaker). It also mixes with ice melt which has lower salinity. Both factors increase water density, so that denser water sinks to the ocean floor, where it is gradually moved by deep ocean currents. This convective channel implies that there is additional heat exchange between the atmosphere and the oceans, in a way that maintains a constant temperature differential between the two.

In this subsection, we show that incorporating this channel into our warming block leaves the quantitative predictions of our climate module virtually identical once properly recalibrated. To do so in a disciplined way, we propose the following modification to the warming module from Geoffroy et al. (2013) that underlies Folini et al. (2024):

$$
\begin{aligned}
C^U \dot{T}_t^U &= F_t - \ell T_t^U + \varepsilon\gamma(T_t^L - T_t^U) + \omega(T_t^L - \phi T_t^U) \\
C^L \dot{T}_t^L &= \gamma(T_t^U - T_t^L) + \omega(\phi T_t^U - T_t^L).
\end{aligned}
$$

C^U, C^L denote the heat capacity of the atmosphere and the ocean, respectively. Importantly, the heat flow $\gamma(T_t^L - T_t^U)$ that enters the atmosphere temperature equation is the same as the one that enters the ocean temperature equation. Geoffroy et al. (2013) note that this heat flow is best adjusted by an additional coefficient ε possibly different from one to incorporate some effects of the thermohaline circulation (Manabe et al., 1991; Winton et al., 2010; Held et al., 2010). However, the adjustment by ε does not resolve the problem that ocean and atmosphere temperatures must equalize in the long run.

Our proposed modification includes an additional term $\omega(T_t^L - \phi T_t^U)$. This adjustment captures the idea that convective exchange between the ocean and the atmosphere depends on polar temperatures in the upper ocean, which are lower than average temperatures in the upper ocean. We then also impose that heat inflows and outflows between the two layers due to thermohaline circulation are identical.

We now discuss how to calibrate the additional parameters ω, ϕ. We start by renormalizing the system into a similar format to the one in Folini et al. (2024):

$$\dot{T}_t^U = n\Big(F_t - \ell T_t^U + b_1(T_t^L - T_t^U)\Big) + c_1(T_t^L - \phi T_t^U)$$
$$\dot{T}_t^L = b_2(T_t^U - T_t^L) + c_2(\phi T_t^U - T_t^L),$$

with $c_1 = \omega/C^U$, $c_2 = \omega/C^L$, $n = 1/C^U$, $\varepsilon\gamma = b_1$ and $b_2 = \gamma/C^L$. The corresponding transition matrix is:

$$\mathbf{B} = \begin{pmatrix} -n(b_1 + \ell) - c_1\phi & nb_1 + c_1 \\ b_2 + c_2\phi & -(b_2 + c_2) \end{pmatrix}.$$

Geoffroy et al. (2013) use $C^U = 8$ W yr m^{-2} K^{-1} and $C^L = 100$ W yr m^{-2} K^{-1}, so that: $c_1/c_2 = 8/100 = 0.08$. Next, we posit that in the initial steady state ϕT^U is approximately 0°C, and T^U is approximately 15°C.

We specify our warming module in Kelvin because it represents the relevant absolute unit in which to measure heat exchange. Kelvin degrees are a constant shift relative to Celsius degrees. Therefore, in the original Geoffroy et al. (2013) and Folini et al. (2024) specifications, switching between Kelvin and Celsius degrees drops out. With our proposed modification, it is isomorphic to a constant shift in exogenous forcing.

We now verify whether this new climate module with a thermohaline circulation adjustment delivers similar predictions to the one without the thermohaline circulation adjustment. Starting from steady-state, we will consider four distinct shocks that mirror the exercises in Folini et al. (2024): one in which the atmospheric carbon content exogenously quadruples at time 0, one in which it exogenously doubles at time 0, one in which the carbon content of the atmosphere grows at 1% per year so that it quadruples after roughly 140 years, and one in which it grows at 3% per year.

Parameters in the baseline warming module have been chosen to best match the temperature paths that follow from these two shocks in large-scale climate models—CMIP5 simulations. Therefore, we also recalibrate the parameters in the module with the thermohaline circulation adjustment to match the same CMIP5 simulations.

We proceed as follows. At any point in time: $\phi = \frac{\text{Polar temperature}}{T^U}$. We use the current level of this ratio to calibrate ϕ. That is: $\phi = \frac{273.15}{288.15} = 0.948$. Let $\iota = \frac{T^L}{T^U}$. We use the pre-industrial steady-state ratio to calibrate: $\iota = \frac{278.15}{288.15} = 0.9653$. In steady state: $(b_2 + c_2\phi)T^U = (b_2 + c_2)T^L$, implying: $\frac{T^L}{T^U} = \frac{b_2+c_2\phi}{b_2+c_2} < 1$. We invert this relationship to obtain c_2: $c_2 = b_2\frac{1-\iota}{\iota-\phi}$ and $c_1 = 0.08 \cdot c_2$.

In our recalibration exercise, we keep n and ℓ at their original values, as they are radiative parameters calibrated on the long-run climate response after the quadrupling of atmospheric carbon concentration, which should be invariant to the model structure. We then choose b_1 and b_2 to minimize the maximum distance between the temperature paths from the baseline Folini et al. (2024) module and our adjusted module in the quadrupling exercise. We use the instantaneous doubling and gradual increases exercises as an over-identification check.

Figure 14 shows the fit of our recalibration exercise. Figures 15, 16 and 17 display the results of our over-identification exercises. Across all exercises, both models behave very similarly.

Figure 14: Quadrupling Carbon Dioxide Concentrations at Time 0

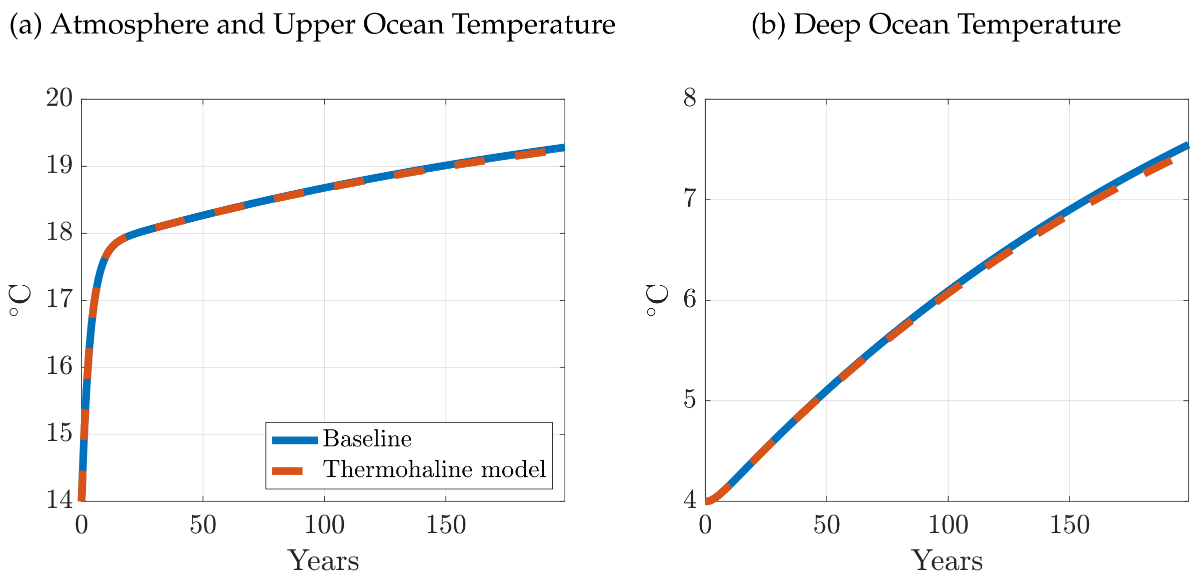

Notes: Response of the climate system to an exogenous quadrupling of atmospheric carbon concentrations at time 0. The figure compares the baseline module from Folini et al. (2024) with the module that includes the thermohaline circulation adjustment, after recalibrating b_1 and b_2. This experiment is the target of the recalibration.

Figure 15: Doubling Carbon Dioxide Concentrations at Time 0

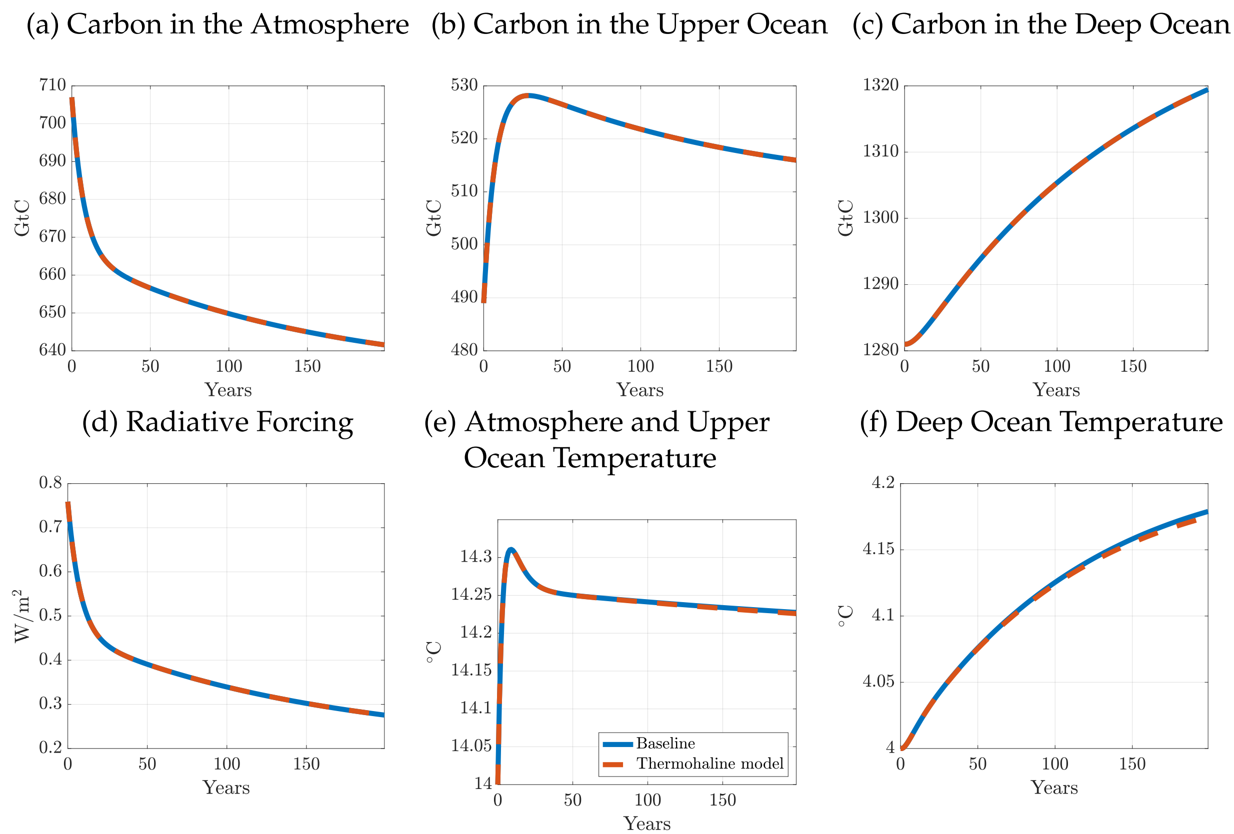

Notes: Response of the climate system to an exogenous doubling of atmospheric carbon concentrations at time 0. The figure compares the baseline module from Folini et al. (2024) with the module that includes the thermohaline circulation adjustment, after recalibrating b_1 and b_2 to the quadrupling experiment in Figure 14. This experiment is an over-identification check.

Figure 16: Constant Increase in Carbon Dioxide Concentrations at 1% a Year

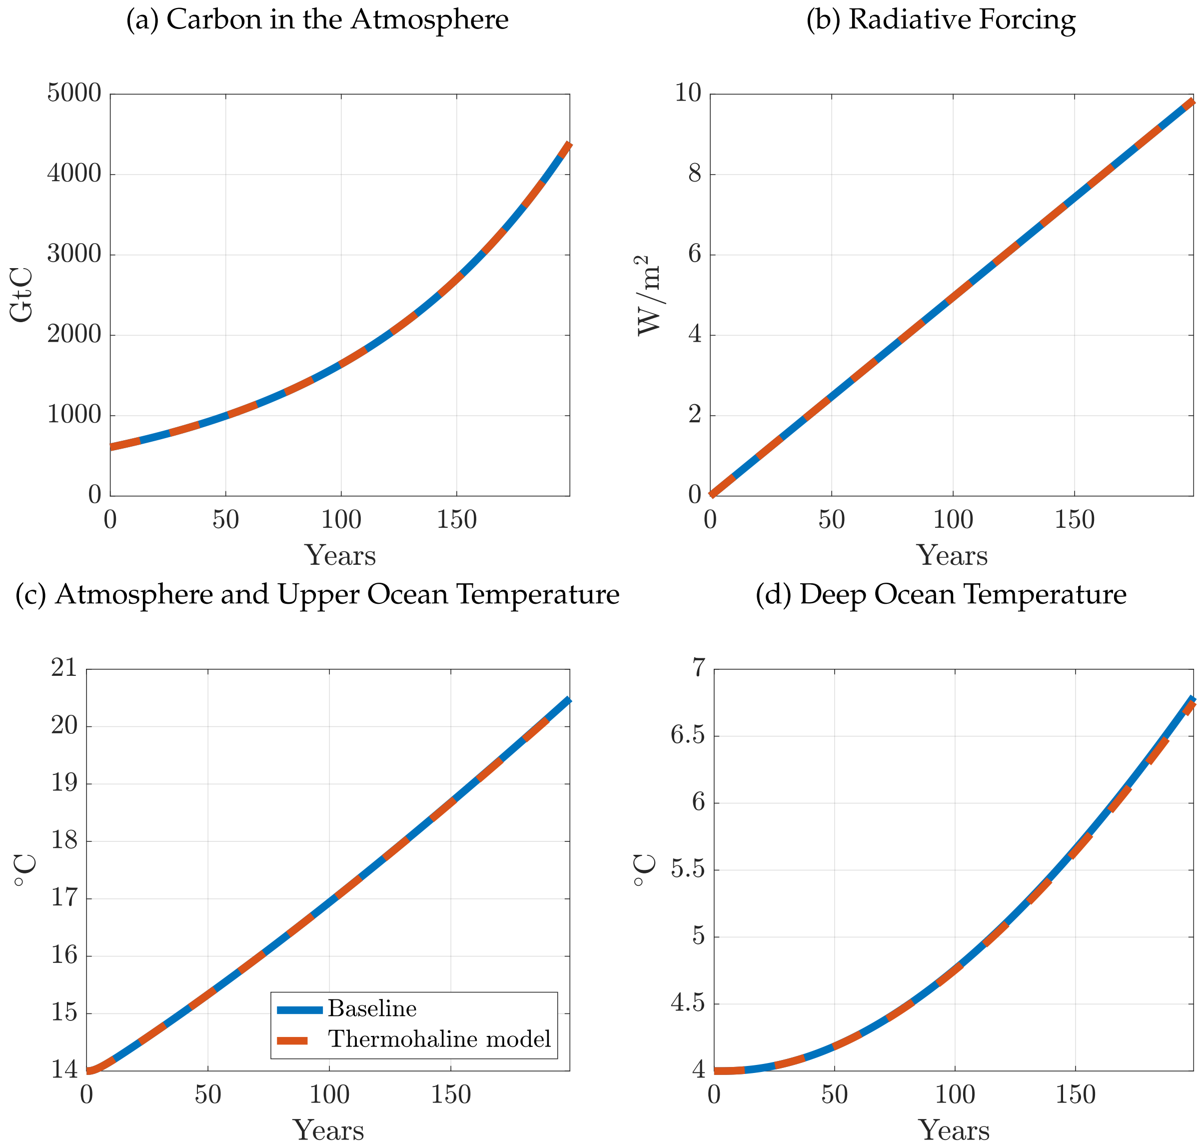

Notes: Response of the climate system when atmospheric carbon concentrations grow exogenously at 1% per year, so that they quadruple after roughly 140 years. The figure compares the baseline module from Folini et al. (2024) with the module that includes the thermohaline circulation adjustment, after recalibrating b_1 and b_2 to the quadrupling experiment in Figure 14. This experiment is an over-identification check.

Figure 17: Constant Increase in Carbon Dioxide Concentrations at 3% a Year

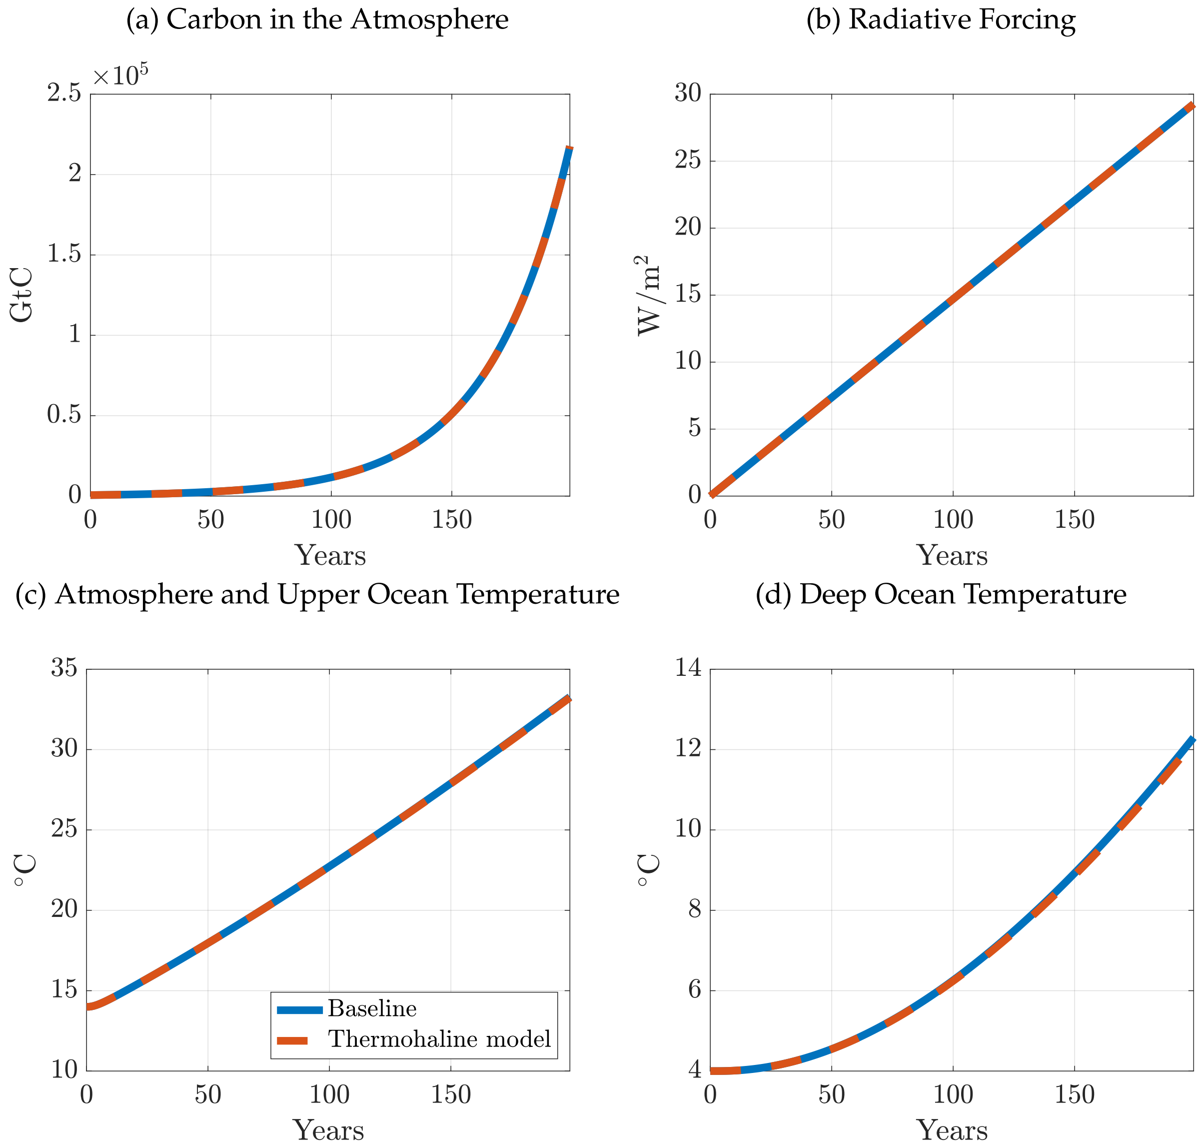

Notes: Response of the climate system when atmospheric carbon concentrations grow exogenously at 3% per year. The figure compares the baseline module from Folini et al. (2024) with the module that includes the thermohaline circulation adjustment, after recalibrating b_1 and b_2 to the quadrupling experiment in Figure 14. This experiment is an over-identification check.

G Numerical Solution

This appendix describes the algorithm that we use to solve for the nonlinear transition paths of the model. We solve for the entire path of all variables at once with a Newton method that uses an exact, sparse, analytic Jacobian. Appendix G.1 summarizes the detrended equilibrium conditions in the laissez-faire and in the planner allocation. Appendix G.2 recasts them as a square system of residual equations. Appendix G.3 describes the Newton update and the analytic Jacobian. Appendix G.4 describes the structure of the solver. Appendix G.5 discusses the terminal steady state of the planner solution and the additional fixed point on the terminal stock of carbon. Appendix G.6 describes the welfare computations. Appendix G.7 describes the variants of the solver used in the backstop and nonlinear damage exercises.

G.1 Detrended System

Detrending. We discretize the model with a period length of one year. All economic variables are expressed in per capita terms and detrended by their balanced growth rates from Propositions 1 and 2. Lowercase letters denote detrended variables. For instance $c_t = C_t e^{-\bar{g}_Y t}$, where $\bar{g}_Y = g_Y - n$ is per capita output growth. Energy is detrended at $g_E = g_Y - g_\chi$ and emissions at $g_m = g_E - g_\sigma$. Population growth is handled by the rescaling of Appendix A.3, so that n enters through the effective discount rate $\rho - n$ and the effective depreciation rate $\delta + n$. When emissions grow, $g_m > 0$, the climate block is detrended as well. The carbon vector is detrended at g_m and both temperature layers share a common linear trend $\theta = hg_m/\ell$, which is removed. Here $h = F_{2XCO2}/\log 2$ denotes the semi-elasticity of forcing to atmospheric carbon, n the inverse thermal capacity of the upper layer, and ℓ radiative losses to outer space, as in Appendix B. With a slight abuse of notation, the detrended carbon and temperature vectors keep their symbols $\mathbf{M}_t$ and $\mathbf{T}_t$.

The damage kernel is discretized into $S = 50$ annual weights ζ_s, together with the $(\rho - n)$-discounted weights ζ_s^d that enter the temperature costate below. In this appendix the kernel carries the sign convention of the code: $\zeta_s \leq 0$ under damages, so that it equals the negative of the kernel in the main text. With this convention the planner costates and the implied carbon tax are positive.

Planner system. The detrended planner allocation solves, for every t,

$$
\begin{aligned}
\left(\frac{c_{t+1}}{c_t}\right)^{\gamma} &= 1 + r_{t+1} - \rho - \gamma \overline{g}_Y \\
y_t &= \left(k_t^{\alpha} (z_t l_t)^{1-\alpha}\right)^{1-\psi} e_t^{\psi} \\
r_t + \delta &= \alpha(1-\psi)\, y_t / k_t \\
\psi\, y_t / e_t &= p_t^e + \sigma_t \tau_t, \qquad \sigma_t \tau_t = \sigma_0 \widetilde{\mu}_{2t}^A c_t^{\gamma} \\
k_{t+1} &= (1 - \delta - g_Y) k_t + y_t - c_t - p_t^e e_t \\
\mathbf{M}_{t+1} &= \left(\mathbf{A} + (1 - g_m)\mathbf{I}\right)\mathbf{M}_t + \sigma_0 e_t\, \mathbf{e}_1 \\
f_t &= h \log\left(M_t^A / M_{eq}^A\right) + F_{eq}^A \\
\mathbf{T}_{t+1} &= (\mathbf{B} + \mathbf{I})\mathbf{T}_t + \mathbf{N}\begin{pmatrix} f_t \\ F_{eq}^L \end{pmatrix} - \theta \begin{pmatrix} 1 \\ 1 \end{pmatrix} \\
\log z_t &= \log z_0 + \sum_{s=1}^{S} \zeta_s \left(T_{t-s}^U - T_{eq}^U\right) - \theta \sum_{s=1}^{S} s\, \zeta_s \\
\widetilde{\boldsymbol{\mu}}_{2,t+1} &= \left((1 + \rho - n + g_m)\mathbf{I} - \mathbf{A}'\right)\widetilde{\boldsymbol{\mu}}_{2t} - \frac{nh\, \mu_{3t}^U}{M_t^A}\, \mathbf{e}_1 \\
\boldsymbol{\mu}_{3,t+1} &= \left((1 + \rho - n)\mathbf{I} - \mathbf{B}'\right)\boldsymbol{\mu}_{3t} + (1-\alpha)(1-\psi) \sum_{s=t+1}^{t+S} \zeta_{s-t}^d \, \frac{y_s}{c_s^{\gamma}}\, \mathbf{e}_1
\end{aligned}
$$

where $\widetilde{\boldsymbol{\mu}}_{2t}$ is the detrended costate on carbon, $\boldsymbol{\mu}_{3t}$ the detrended costate on temperature, and τ_t the carbon tax from equation (11). Investment has been substituted out through the resource constraint. The constant $\theta \sum_s s\zeta_s$ in the productivity equation is the residue of removing the temperature trend inside the damage kernel. In the interior regime the detrended energy price is constant, $p_t^e = \chi_0$. In the capped-emissions regime, energy is instead detrended at $g_E = \overline{g}_m + g_\sigma$, and the detrended extraction cost satisfies $p_t^e = \chi_0 e^{(g_\chi + g_E - g_Y)t} \to 0$, because the tax dominates extraction costs in the long run.

Laissez-faire system. The laissez-faire satisfies the same equations without the costates and with $\tau_t = 0$, so that energy demand reads $\psi y_t / e_t = p_t^e = \chi_0$.

G.2 Residual System

We recast the discretized system as a square system of residual equations. Script letters denote the residual blocks in this appendix. Rearranging each equation above so that it equals zero at a solution yields, in the planner case, the Euler residuals $\mathcal{E}_t$, the rental rate residuals $\mathcal{R}_t$, the production residuals $\mathcal{Y}_t$ (in logs), the energy pricing residuals $\mathcal{P}_t$ (with the tax substituted in), the productivity residuals $\mathcal{Z}_t$ (in logs), the capital accumulation residuals $\mathcal{G}_t$, the carbon cycle residuals $\mathcal{M}_t$, the temperature residuals $\mathcal{T}_t$ (with forcing substituted in), the temperature costate residuals $\mathcal{W}_t$, and the carbon costate residuals $\mathcal{S}_t$. The blocks $\mathcal{M}_t, \mathcal{S}_t$ are three-dimensional and $\mathcal{T}_t, \mathcal{W}_t$ are two-dimensional. Taking the production and productivity equations in logs makes those rows exactly linear in logs and in temperatures.

Boundary conditions. The economy starts on the calibrated laissez-faire path in 2020. This provides the initial conditions k_0, $\mathbf{M}_0$, $\mathbf{T}_0$, together with the pre-period upper layer temperatures that enter the damage kernel during the first S periods. We approximate the infinite horizon with a finite path of length T, whose last period is the steady state of the detrended system. All period-T variables are fixed at their steady-state values. Any lead beyond T in the costate sums is evaluated at steady-state values as well. We use $T = 1400$ for the laissez-faire and $T = 1500$ for the planner.

Counting. The planner unknowns are c_t, r_t, y_t, e_t, z_t for $t = 0, \ldots, T-1$, the states k_t, $\mathbf{M}_t$, $\mathbf{T}_t$ for $t = 1, \ldots, T-1$, and the costates $\boldsymbol{\mu}_{3t}$, $\widetilde{\boldsymbol{\mu}}_{2t}$ for $t = 0, \ldots, T-1$, for a total of $16T - 6$ scalars. The residual equations count $16T - 6$ as well, so the system is square. Three details matter for the counting. First, there are no $\mathcal{G}_{T-1}$, $\mathcal{M}_{T-1}$, $\mathcal{T}_{T-1}$ equations, which would over-determine the terminal states. Second, the last Euler and costate equations enforce the terminal values. Third, initial conditions enter the date-0 equations as constants. The laissez-faire system drops the costate blocks and the tax, leaving $11T - 6$ unknowns and equations.

G.3 Newton Update and Analytic Jacobian

Stack all unknowns in a vector $\mathbf{X}$ and all residuals in a map $\mathcal{H}(\mathbf{X})$. The Newton update is

$$\mathbf{X}^{n+1} = \mathbf{X}^n - \varsigma_n \, \mathcal{J}(\mathbf{X}^n)^{-1} \mathcal{H}(\mathbf{X}^n),$$

where $\mathcal{J}$ is the Jacobian of $\mathcal{H}$ and $\varsigma_n \in (0, 1]$ is a backtracking line search parameter that we shrink until the update reduces $\|\mathcal{H}\|$. Near the solution $\varsigma_n = 1$ and convergence is quadratic.

Positivity. Far from the solution, the full Newton step in levels can push positive variables negative. We therefore update the positive variables $(c_t, y_t, e_t, z_t, k_t, \mathbf{M}_t)$ multiplicatively. We rescale their columns of $\mathcal{J}$ into log units, solve for the step d, and update $x \leftarrow x e^{\varsigma_n d}$. Variables that can change sign, namely r_t, $\mathbf{T}_t$ and the costates, are updated in levels. The log step coincides with the level step to first order, so quadratic convergence is unaffected, and every iterate remains positive.

Sparsity. The Jacobian is sparse. Each equation dated t involves a handful of variables dated t and $t+1$, plus S lags of temperature in $\mathcal{Z}_t$ and S leads of output and consumption in $\mathcal{W}_t$. The table below lists the time offsets j such that $\partial \mathcal{X}_t / \partial v_{t+j} \neq 0$ for each equation block $\mathcal{X}$ and variable block v. A superscript restricts the entry to the atmospheric carbon component (A) or the upper temperature layer component (U).

	c	r	y	e	z	k	$\mathbf{M}$	$\mathbf{T}$	$\boldsymbol{\mu}_3$	$\widetilde{\boldsymbol{\mu}}_2$
$\mathcal{E}_t$	$0, +1$	$+1$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$
$\mathcal{R}_t$	$\cdot$	0	0	$\cdot$	$\cdot$	0	$\cdot$	$\cdot$	$\cdot$	$\cdot$
$\mathcal{Y}_t$	$\cdot$	$\cdot$	0	0	0	0	$\cdot$	$\cdot$	$\cdot$	$\cdot$
$\mathcal{P}_t$	0	$\cdot$	0	0	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	0^A
$\mathcal{Z}_t$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	0	$\cdot$	$\cdot$	$-S, \ldots, -1$ (U)	$\cdot$	$\cdot$
$\mathcal{G}_t$	0	$\cdot$	0	0	$\cdot$	$0, +1$	$\cdot$	$\cdot$	$\cdot$	$\cdot$
$\mathcal{M}_t$	$\cdot$	$\cdot$	$\cdot$	0	$\cdot$	$\cdot$	$0, +1$	$\cdot$	$\cdot$	$\cdot$
$\mathcal{T}_t$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	0^A	$0, +1$	$\cdot$	$\cdot$
$\mathcal{W}_t$	$+1, \ldots, +S$	$\cdot$	$+1, \ldots, +S$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$0, +1$	$\cdot$
$\mathcal{S}_t$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	0^A	$\cdot$	0^U	$0, +1$

Jacobian entries. Each nonzero entry is an elementary derivative of the residuals above. Writing $\kappa_w = (1-\alpha)(1-\psi)$, the entries are

$$
\begin{aligned}
&\frac{\partial \mathcal{E}_t}{\partial c_t} = -\frac{\gamma}{c_t}\Big(\frac{c_{t+1}}{c_t}\Big)^{\gamma} && \frac{\partial \mathcal{E}_t}{\partial c_{t+1}} = \frac{\gamma}{c_{t+1}}\Big(\frac{c_{t+1}}{c_t}\Big)^{\gamma} && \frac{\partial \mathcal{E}_t}{\partial r_{t+1}} = -1 \\
&\frac{\partial \mathcal{R}_t}{\partial r_t} = 1 && \frac{\partial \mathcal{R}_t}{\partial y_t} = -\frac{\alpha(1-\psi)}{k_t} && \frac{\partial \mathcal{R}_t}{\partial k_t} = \alpha(1-\psi)\frac{y_t}{k_t^2} \\
&\frac{\partial \mathcal{Y}_t}{\partial y_t} = \frac{1}{y_t} && \frac{\partial \mathcal{Y}_t}{\partial k_t} = -\frac{\alpha(1-\psi)}{k_t} && \frac{\partial \mathcal{Y}_t}{\partial e_t} = -\frac{\psi}{e_t} \\
&\frac{\partial \mathcal{Y}_t}{\partial z_t} = -\frac{\kappa_w}{z_t} && \frac{\partial \mathcal{P}_t}{\partial \widetilde{\mu}^A_{2t}} = \sigma_0 c_t^{\gamma} && \frac{\partial \mathcal{P}_t}{\partial c_t} = \gamma\sigma_0\widetilde{\mu}^A_{2t}c_t^{\gamma-1} \\
&\frac{\partial \mathcal{P}_t}{\partial y_t} = -\frac{\psi}{e_t} && \frac{\partial \mathcal{P}_t}{\partial e_t} = \psi\frac{y_t}{e_t^2} && \frac{\partial \mathcal{Z}_t}{\partial z_t} = \frac{1}{z_t} \\
&\frac{\partial \mathcal{Z}_t}{\partial T^U_{t-s}} = -\zeta_s && \frac{\partial \mathcal{G}_t}{\partial k_{t+1}} = 1 && \frac{\partial \mathcal{G}_t}{\partial k_t} = -(1-\delta-g_Y) \\
&\frac{\partial \mathcal{G}_t}{\partial y_t} = -1 && \frac{\partial \mathcal{G}_t}{\partial c_t} = 1 && \frac{\partial \mathcal{G}_t}{\partial e_t} = p^e_t
\end{aligned}
$$

for the economic blocks, with $s = 1, \ldots, S$ in the $\mathcal{Z}_t$ rows, and

$$
\begin{aligned}
&\frac{\partial \mathcal{M}_t}{\partial \mathbf{M}_{t+1}} = \mathbf{I} && \frac{\partial \mathcal{M}_t}{\partial \mathbf{M}_t} = -\big(\mathbf{A} + (1-g_m)\mathbf{I}\big) && \frac{\partial \mathcal{M}_t}{\partial e_t} = -\sigma_0\mathbf{e}_1 \\
&\frac{\partial \mathcal{T}_t}{\partial \mathbf{T}_{t+1}} = \mathbf{I} && \frac{\partial \mathcal{T}_t}{\partial \mathbf{T}_t} = -(\mathbf{B}+\mathbf{I}) && \frac{\partial \mathcal{T}_t}{\partial M^A_t} = -\frac{h}{M^A_t}\mathbf{N}\mathbf{e}_1 \\
&\frac{\partial \mathcal{W}_t}{\partial \boldsymbol{\mu}_{3t+1}} = \mathbf{I} && \frac{\partial \mathcal{W}_t}{\partial \boldsymbol{\mu}_{3t}} = -\big((1+\rho-n)\mathbf{I} - \mathbf{B}'\big) && \frac{\partial \mathcal{W}_t}{\partial y_s} = -\kappa_w\frac{\zeta^d_{s-t}}{c_s^{\gamma}}\mathbf{e}_1 \\
&\frac{\partial \mathcal{W}_t}{\partial c_s} = \gamma\kappa_w\frac{y_s\zeta^d_{s-t}}{c_s^{\gamma+1}}\mathbf{e}_1 && \frac{\partial \mathcal{S}_t}{\partial \widetilde{\boldsymbol{\mu}}_{2t+1}} = \mathbf{I} && \frac{\partial \mathcal{S}_t}{\partial \widetilde{\boldsymbol{\mu}}_{2t}} = -\big((1+\rho-n+g_m)\mathbf{I} - \mathbf{A}'\big) \\
&\frac{\partial \mathcal{S}_t}{\partial \mu^U_{3t}} = \frac{nh}{M^A_t}\mathbf{e}_1 && \frac{\partial \mathcal{S}_t}{\partial M^A_t} = -\frac{nh\,\mu^U_{3t}}{(M^A_t)^2}\mathbf{e}_1
\end{aligned}
$$

for the climate and costate blocks, with $s = t+1, \ldots, t+S$ in the $\mathcal{W}_t$ rows. Entries whose column corresponds to a boundary value rather than an unknown are simply absent. With a period-interleaved ordering of the unknowns, $\mathcal{J}$ is a banded matrix with a half-bandwidth of about S periods. Sparse LU factorization then makes each Newton step linear in the horizon T, in time and in memory.

G.4 Structure of the Solver

Initial point. The initial point comes from one block elimination pass of the system. Given guessed sequences for the interest rate and energy use (r_t, e_t), we solve consumption backward from the Euler equations, the climate block and productivity forward from energy use, capital and output from the static equations, and the costates backward from their terminal values, which delivers the tax. At the resulting point every residual is zero except the resource constraint rows $\mathcal{G}$ and the energy pricing rows $\mathcal{P}$. These two sequences are the natural residuals of the underlying fixed point problem in (r_t, e_t). The Newton method on the full stacked system is exactly a block elimination of this reduced problem, with the advantage that the stacked Jacobian remains sparse whereas the reduced one is dense. Guesses come from simple design paths for the baseline solves, and from continuation along the damage grid for the welfare curves of Figure 6.

Convergence. From these initial points the laissez-faire converges in about 6 Newton iterations and the planner in about 9, to residuals at machine precision. A full planner solve at $T = 1500$ takes about one second. On rare occasions near the regime change threshold the line search stalls, in which case the solver finishes with a quasi-Newton step using the same analytic Jacobian.

G.5 Terminal Steady State of the Planner

In the interior regime the detrended planner steady state is unique, and the terminal values are genuine constants of the residual system. In the capped-emissions regime it is not. Extraction costs become negligible relative to the carbon tax, and there is a one-dimensional continuum of detrended steady states, indexed by the terminal carbon stock. The correct boundary condition is a self-consistency requirement. The path must converge to the steady state whose carbon stock is the one the path itself accumulates,

$$\text{terminal values} = \overline{x}(\varkappa), \qquad \varkappa = \mathbf{1}'\mathbf{M}_T,$$

where $\overline{x}(\varkappa)$ collects the steady-state values that enter the residual system.

This requirement adds a fixed point on the terminal stock of carbon, which the solver handles in two places. First, every evaluation of the residuals recomputes the terminal steady state from the carbon stock that the current path accumulates. The elimination pass described above therefore loops on the terminal carbon stock automatically. Second, the Newton Jacobian must include the dependence of the residuals on the unknowns through the terminal values. We promote $\mathbf{M}_T$ to an unknown, restore the carbon cycle equation $\mathcal{M}_{T-1}$, and apply the chain rule to the composite residual map. The correction is a rank-one block confined to the columns of $\mathbf{M}_T$. It affects the last Euler equation through $\bar{c}$, the post-horizon terms of the temperature costate through $\bar{y}/\bar{c}$, and the last costate equations through $\bar{\boldsymbol{\mu}}_3$ and $\bar{\bar{\boldsymbol{\mu}}}_2$. Because the steady-state interest rate is the modified golden rule, it does not depend on $\varkappa$. Because consumption is a fixed share of output in this regime, $\bar{y}/\bar{c}$ and $\bar{\boldsymbol{\mu}}_3$ are invariant to $\varkappa$ as well, and their entries vanish. The remaining sensitivities $\bar{x}'(\varkappa)$ are computed with one scalar finite difference of the steady-state computation per Jacobian assembly, which is essentially free.

Omitting the correction while still refreshing the terminal values yields an inexact Newton method. Convergence near the solution then degrades from quadratic to linear. At the baseline calibration the degradation is mild, 12 instead of 9 iterations, but the surviving elasticities are of order one, so we always include the correction.

G.6 Welfare

Welfare comparisons use the exact discounted sum of log consumption implied by each solved path. For an allocation a with detrended consumption path c_t^a and trend growth $\bar{g}_Y^a$, welfare under log utility and effective discounting $\rho - n$ is

$$W_a = \sum_{t=0}^{T-1} \beta^t \big(\log c_t^a + \bar{g}_Y^a t \big) + \beta^T \left[\frac{\log \bar{c}^{\,a} + \bar{g}_Y^a T}{1-\beta} + \frac{\bar{g}_Y^a \beta}{(1-\beta)^2} \right], \qquad \beta = e^{-(\rho - n)},$$

where the second term is the closed-form continuation from the terminal steady state. Every allocation is integrated over its own horizon, with its own trend and its own terminal consumption level. The consumption-equivalent variation relative to the laissez-faire is $\text{CEV} = \exp\big((\rho - n)(W_a - W_{lf})\big) - 1$.

G.7 Variants

Constant temperature. The counterfactual in which temperature stops increasing in 2020 switches off the climate block entirely. Productivity becomes an exogenous sequence and the residual system reduces to the economic blocks in $(c_t, r_t, y_t, e_t, k_t)$. The Jacobian is banded with a one-period bandwidth, and the same solver converges in about 4 iterations.

Backstop technology. The backstop exercise of Figure 7 holds temperature constant from the arrival date H onward, at the level reached at H. Before H the system is the baseline one. After H the climate block drops out, and detrended variables use the constant-temperature growth rates, spliced continuously at the seam. Post-backstop productivity loads the frozen temperature T^U_H through the cumulative kernel weights, so the corresponding Jacobian rows combine derivatives from the two regimes. The terminal steady state is endogenous through T^U_H alone, with analytic sensitivities. In the planner solution the backstop changes the boundary conditions of the costates and freezes the detrended energy price at its seam value from H onward. The carbon costate satisfies $\widetilde{\boldsymbol{\mu}}_{2H} = 0$, because emissions dated H or later never affect temperature, so the optimal tax declines to zero at the backstop date. The temperature costate at H prices the stabilized temperature through the discounted cumulative kernel over all post-backstop dates, with a closed-form tail.

Nonlinear damages. The nonlinear damage exercise of Figure 8 replaces the productivity equation with $\log z_t = \log z_1 - (D(T^U_t) - D(T^U_1))$ for a general damage function D anchored at the pre-industrial temperature, with damages contemporaneous in the true temperature. Because no balanced growth path exists under nonlinear damages and trending temperature, we always impose a backstop at $H = 1000$ and detrend with the constant-temperature rates in both regimes. Temperatures then remain in levels, and the detrended energy price is constant for the laissez-faire and for the planner. The kernel band of the Jacobian collapses to a diagonal loading $D'(T^U_t)$, the costate source becomes the own-period marginal damage $(y_t/c_t)D'(T^U_t)$, and the temperature costate rows gain a curvature entry proportional to $D''(T^U_t)$. The terminal steady state is the constant-temperature steady state indexed by T^U_H, with exact analytic sensitivities as in the backstop exercise. Because the detrended carbon stock of a mitigating path decays through roughly thirteen orders of magnitude over such a long horizon, the solver carries the carbon block and its costate in true units, an exact reparametrization that keeps all Jacobian entries bounded. The line search accepts a step when it reduces the norm of the Newton-transformed residual, which is insensitive to the large tax and costate rows near H. The laissez-faire solution under strongly nonlinear damages is seeded by a damped fixed point on the climate and economic blocks, started at the planner solution.

References Appendix

Aghion, Philippe and Peter Howitt (1992). “A Model of Growth through Creative Destruction”. *Econometrica* 60.2, pp. 323–351.

Bilal, Adrien and Diego R Känzig (2026). “The Macroeconomic Impact of Climate Change: Global Versus Local Temperature”. *The Quarterly Journal of Economics* 141.2, pp. 889–944.

Folini, Doris, Aleksandra Friedl, Felix Kübler, and Simon Scheidegger (2024). “The Climate in Climate Economics”. *The Review of Economic Studies*.

Geoffroy, Olivier, David Saint-Martin, Dirk J.L. Olivié, Aurore Voldoire, Gilles Bellon, and Sophie Tytéca (2013). “Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments”. *Journal of climate* 26.6, pp. 1841–1857.

Grossman, Gene M and Elhanan Helpman (1991). “Trade, knowledge spillovers, and growth”. *European Economic Review* 35.2-3, pp. 517–526.

Held, Isaac M., Michael Winton, Ken Takahashi, Thomas Delworth, Fanrong Zeng, and Geoffrey K. Vallis (2010). “Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing”. *Journal of Climate* 23.9, pp. 2418–2427.

Manabe, Syukaro, Ronald J. Stouffer, Michael J. Spelman, and Kirk Bryan (1991). “Transient responses of a coupled ocean–atmosphere model to gradual changes of atmospheric CO2. Part I. Annual mean response”. *Journal of Climate* 4.8, pp. 785–818.

NOAA National Centers for Environmental Information (2023a). *Climate at a Glance: Global Time Series*. Online Resource.

— (2023b). *Global Surface Temperature Anomalies: Mean Temperature Estimates*. Online Resource.

Winton, Michael, Ken Takahashi, and Isaac M. Held (2010). “Importance of ocean heat uptake efficacy to transient climate change”. *Journal of Climate* 23.9, pp. 2333–2344.