Grantham Research Institute on Climate Change and the Environment

Optimal climate policy as if the transition matters

Emanuele Campiglio, Simon Dietz and Frank Venmans

May 2026

Grantham Research Institute on Climate Change and the Environment
Working Paper No. 387

ISSN 2515-5717 (Online)

Grantham Foundation for the Protection of the Environment

LSE Global School of Sustainability

The Grantham Research Institute on Climate Change and the Environment was established in 2008 at the London School of Economics and Political Science. The Institute brings together international expertise on economics, as well as finance, geography, the environment, international development and political economy to establish a world-leading centre for policy-relevant research, teaching and training in climate change and the environment. It is hosted by the Global School of Sustainability at LSE and is funded by the Grantham Foundation for the Protection of the Environment.

www.lse.ac.uk/granthaminstitute

This paper was first published in May 2026 by the Grantham Research Institute on Climate Change and the Environment at the London School of Economics and Political Science.

Suggested citation:

Campiglio E, Dietz S and Venmans F (2026) *Optimal climate policy as if the transition matters*. Grantham Research Institute on Climate Change and the Environment Working Paper 387. London: London School of Economics and Political Science

Optimal climate policy as if the transition matters*

Emanuele Campiglio[a,b,c], Simon Dietz[c], and Frank Venmans[c]

[a] *University of Bologna, Department of Economics, Bologna, Italy*
[b] *RFF-CMCC European Institute on Economics and the Environment (EIEE), Milan, Italy*
[c] *London School of Economics and Political Science, London, UK*

May 7, 2026

Abstract

We study optimal climate policy when the low-carbon transition is constrained by frictions in scaling clean energy and retiring fossil assets. We develop a stochastic, quantitative, intertemporal general-equilibrium model with heterogeneous dirty capital disciplined using firm-level data and adjustment costs disciplined using evidence on historical energy-system expansions and phase-outs. The data reveal an important 'dirty tail' of highly emissions-intensive assets that can be retired early at low cost. This means that, despite transition frictions, optimal policy entails rapid near-term decarbonisation, with emissions halving in the first ten years. The transition is not primarily constrained by the cost of stranding dirty capital, but by the rate at which clean capital can be expanded. Our results highlight the importance of modelling heterogeneity in emissions intensity and support policies targeting high-emissions assets and easing clean investment constraints.

Keywords: adjustment costs; carbon price; climate change; low-carbon transition; stranded assets; technological progress; uncertainty
JEL codes: C61, E22, H23, O44, Q54, Q55

*Email for correspondence: F.Venmans1@lse.ac.uk. We would like to thank the editor, three anonymous reviewers, Ghassane Benmir, Yongyang Cai, Derek Lemoine, Rick van der Ploeg and Ivan Rudik for comments, as well as participants at various seminars, workshops and conferences. We would also like to thank Simon Scheidegger for guiding us through his neural nets code.

1 Introduction

Reducing greenhouse gas emissions is becoming more and more urgent. The 2015 Paris Agreement on climate change aimed to hold the increase in global average temperature to well below 2 °C above pre-industrial levels, while pursuing efforts to limit warming to 1.5 °C. These were ambitious goals from the outset. Yet, since the Agreement was signed, global emissions have continued to rise. As a result, any pathway that aims to meet these goals now necessarily involves steep emissions reductions starting now.

While the Paris goals are the outcome of a political process, economic research increasingly reinforces the case for rapid mitigation. Recent estimates of the social cost of carbon (SCC) – the marginal damage caused by an additional tonne of CO_2 – are much higher than most earlier modelling. For example, the U.S. Environmental Protection Agency has leveraged recent, fine-grained econometric work to put the SCC at US\$243/tCO_2 (EPA, 2023), Moore et al. (2024) have used an expert survey of SCC modellers to produce estimates in the range \$169-336/tCO_2, and Bilal and Känzig (2026) estimate an SCC exceeding \$1,200/tCO_2 based on an econometric approach exploiting global rather than local temperature shocks.[1] With such large marginal benefits, rapid near-term emissions reductions are desirable.

At the same time, concerns about the feasibility and cost of rapid decarbonisation remain. Achieving large emissions reductions requires the rapid scaling of clean energy, as well as the rapid phase-out of dirty energy (IEA, 2021). In practice, both margins are subject to frictions. Expanding clean energy is limited by manufacturing capacity, supply chains, permitting and siting processes, etc. Similarly, early retirement of fossil assets entails not only economic losses but also legal, institutional, and political costs (van der Ploeg and Rezai, 2020b). These transition frictions raise the concern that aggressive climate policy may impose large adjustment costs, potentially altering the optimal timing and composition of emissions reductions. The central economic question, therefore, is not whether emissions should be reduced, but how rapidly the transition should proceed once the costs and constraints of adjustment are taken into account.

Accordingly, the core contribution of this paper is to quantify optimal climate policy when the low-carbon transition is constrained by frictions to capital reallocation and the early retirement of fossil assets, within a fully intertemporal, general-equilibrium framework. A central feature of our approach is the explicit modelling of heterogeneous dirty capital. Emissions arise from the use of fossil assets that differ in emissions intensity, and we allow the planner to retire these assets prematurely at a cost. We discipline the composition of dirty capital using global firm- and plant-level data, constructing an empirically grounded

[1] All values are inflated to 2024 prices.

mapping from asset values to emissions. This leads to the crucial insight that *large emissions reductions can be achieved through the early retirement of a relatively small subset of highly emissions-intensive assets, which we call the 'dirty tail'.*

We embed this representation of dirty capital in a model with convex adjustment costs that penalise rapid changes in capital stocks. These costs capture frictions in scaling up clean technologies and in managing large-scale asset retirement. The parameters governing adjustment costs are disciplined using historical evidence on rapid energy-system expansions and phase-outs. Coupled with limited short-run substitutability of dirty energy in production, our approach avoids overstating the ease of transition and ensures that the speed of decarbonisation is endogenously determined.

Because technical change and uncertainty are also important for the transition, we incorporate both. Clean technical change occurs exogenously, and endogenously through learning-by-doing. We introduce uncertainty through a stochastic extension of the model with recursive preferences, allowing us to study how risk affects optimal carbon pricing and the speed of the transition. We incorporate multiple sources of uncertainty, including the business cycle and climate variability. This results in a stochastic model with six states, including four capital stocks, and three control variables, putting it with frontier stochastic Integrated Assessment Models (IAMs). We solve it using state-of-the-art methods based on neural nets.

Our central result is that, even when empirically grounded transition frictions are taken seriously, optimal climate policy entails rapid near-term decarbonisation. The key mechanism is the dirty tail of very emissions-intensive assets, the early retirement of which achieves large emissions reductions at low cost. As a result, the optimal transition begins with a sharp initial decline in emissions, followed by a more gradual reconfiguration of the capital stock as clean capital expands and remaining dirty capital is replaced. This highlights the importance of modelling heterogeneity in emissions intensity, which is central to understanding the speed and structure of the transition.

A second key insight is that transition frictions operate asymmetrically across margins. The transition is not primarily constrained by the cost of stranding dirty capital, but by the rate at which clean capital can be expanded. While the planner does retire emissions-intensive assets, the associated stranding costs are not strongly binding in equilibrium. In contrast, clean investment operates against a tight adjustment constraint, making the expansion of clean capacity the central bottleneck for decarbonisation.

Extending the model to a stochastic setting reinforces these conclusions: optimal investment and emissions are comparatively insensitive to shocks, with uncertainty primarily affecting temperature and the SCC. Early in the transition, emissions reductions are effectively pinned down by the optimal retirement of the dirty tail, making the initial path largely

invariant to realised shocks. Later, emissions reductions are driven by clean-capital expansion and substitution away from remaining dirty capital, but adjustment costs and the high cost of substituting away from remaining dirty capital again make the path insensitive to shocks. By contrast, temperature remains highly uncertain for similar emissions paths, and with convex damages this generates substantial variation in marginal damages/the SCC.

These results highlight the potential importance of policies that directly target the dirty tail, particularly in a second-best setting where it is infeasible to set a globally harmonised carbon price. Examples include targeted finance to support the early retirement of high-emissions assets in developing countries. They also highlight the importance of policies that relax bottlenecks to clean capital investment, given the pace of the transition is governed by constraints on clean capital expansion rather than the cost of stranding dirty assets.

Related literature

Our analysis is situated at the intersection of several strands of climate-economics literature that differ in their modelling objectives, time horizons, and treatment of transition dynamics.

A foundation for much of the relevant literature is given by long-run IAMs designed to study optimal climate policy over centuries, pioneered by William Nordhaus (e.g. Nordhaus, 1991; Barrage and Nordhaus, 2024). These couple models of long-run economic growth with simple climate models. Several strands of literature derive from this, including analytical models that generate closed-form expressions – policy rules – for the optimal carbon tax under simplifying assumptions (e.g. Golosov et al., 2014; Gerlagh and Liski, 2018; Traeger, 2023), models with Directed Technical Change (DTC), in which the transition is driven by endogenous innovation and technology choice (e.g. Acemoglu et al., 2012), and stochastic IAMs that compute an adaptive optimal path under uncertainty (e.g. Lemoine and Traeger, 2014; Cai, 2019; Barnett et al., 2020). Taken together, these approaches are widely used to inform discussions about optimal long-run temperatures, the urgency of emissions reductions, and the SCC. However, transition constraints are typically not the organising principle of the analysis. Transition dynamics are instead simplified either by assuming frictionless reallocation of production factors across uses or frictionless adjustment of an emissions control variable, rather than being the outcome of capital turnover (Hambel et al., 2024, is an exception).

Distinct from long-run IAMs, short-run macroeconomic models have also been applied to climate change in order to study the interaction between climate policy and macroeconomic volatility (see Annicchiarico et al., 2022, for a literature review). Dynamic stochastic general equilibrium (DSGE) models are the core tools used to study these interactions. Applications of these models to climate change are usually based on perturbations around a

steady state, often for individual economies or regions, and are not designed to characterise long-run optimal climate policy or the SCC. Capital is often omitted. Thus, while recent contributions move beyond local dynamics and allow for non-stationary transition paths (Sahuc et al., 2024), their primary focus remains on macroeconomic volatility and monetary policy interactions driven by nominal and financial frictions, rather than on constrained capital reallocation during the low-carbon transition.

Finally, a growing literature studies stranded assets directly. The notion of stranded capital assets originates from empirical studies calculating 'committed emissions' from existing carbon-intensive infrastructure (e.g. Tong et al., 2019). Achieving stringent temperature targets is likely to trigger premature retirement and capital losses from dirty assets, especially in the power sector (e.g. Fofrich Navarro et al., 2026). Other contributions use more stylised representations of dirty capital, analysing how capacity constraints, adjustment costs, and policy design shape the sequencing of investment and capital utilisation (e.g. Baldwin et al., 2020; Rozenberg et al., 2020). We build on this literature by embedding capital accumulation and premature retirement in a dynamic general-equilibrium framework with heterogeneous dirty technologies, allowing the optimal transition path to be shaped by the selective retirement of emissions-intensive assets, rather than by sector-specific constraints or aggregate representations of capital.

The remainder of the paper is organised as follows. Section 2 presents the model. Section 3 analyses how adjustment costs and asset stranding affect the optimal carbon price in theory. Both deterministic and stochastic settings are discussed. Section 4 describes how our quantitative model is calibrated, with a focus on heterogeneous dirty capital, and adjustment costs. Sections 5 and 6 present quantitative results for the deterministic and stochastic models, respectively. Section 7 provides a discussion.

2 The model

In this section, we present a continuous-time version of the model, which is more convenient for deriving analytical solutions. We use a discrete-time version of the model for our numerical simulations.

Production of final goods

A final good Y is used for consumption and investment. The production function is Cobb-Douglas in labour and a composite of capital and energy. Capital and energy are combined

through a constant-elasticity-of-substitution (CES) aggregator:

$$Y(t) = [A(t)L(t)]^{1-\alpha} \left[\varphi_f K_f(t)^{\epsilon_f} + \varphi_E E(t)^{\epsilon_f}\right]^{\alpha/\epsilon_f} \Omega(t), \tag{1}$$

where A is labour-augmenting productivity, L is labour, K_f is capital used in final-good production and E is energy. The parameter α is the share of the capital-energy composite in production, while φ_f and φ_E are CES share parameters satisfying $\varphi_f + \varphi_E = 1$. The parameter $\epsilon_f \equiv (\sigma_f - 1)/\sigma_f$, where σ_f denotes the elasticity of substitution between final-good capital and energy. The multiplier Ω represents climate damages and is explained below. We assume TFP grows exogenously at rate g and labour at rate g_L.

Energy

Energy is produced using two types of capital: dirty and clean. Dirty energy represents fossil fuels without carbon capture and storage (CCS). Clean energy comprises renewables (solar, wind, hydro and modern biofuels), nuclear power, and fossil fuels with CCS. The two energy types are also combined through a CES aggregator,

$$E(t) = \left[\varphi_d (K_d(t) + \xi)^{\epsilon(t)} + \varphi_c K_c(t)^{\epsilon(t)}\right]^{1/\epsilon(t)}, \tag{2}$$

where K_d and K_c denote dirty and clean capital, respectively. The parameters φ_d and φ_c capture their shares and satisfy $\varphi_d + \varphi_c = 1$.[2]

While a nested CES structure is a popular and natural way to represent substitution between clean and dirty energy, an unshifted CES aggregator exhibits a drawback that we wish to highlight: as the dirty-energy input approaches zero, its marginal product becomes unbounded. This implies an unrealistic infinite marginal abatement cost at full decarbonisation. To avoid this behaviour, we include a small positive shifter ξ on dirty capital.

The parameter $\epsilon(t) \equiv (\sigma(t) - 1)/\sigma(t)$, where σ denotes the elasticity of substitution between dirty and clean energy and is allowed to vary over time and with the state of the energy system; its full specification, which captures both exogenous and endogenous technical change, is introduced below. An increasing elasticity of substitution is in line with research showing that the elasticity in the short run is lower than in the long run (e.g. Hassler

[2]In principle, energy production uses both capital and labour, and skills shortages may contribute to transition frictions. However, the labour share of energy value added is empirically small (Barrage, 2020; Fried, 2018). Introducing labour to clean- and dirty-energy production would also require solving a high-dimensional labour-allocation problem at each instant in time. To preserve computational tractability while maintaining appropriate aggregate factor shares, we incorporate all labour inputs into aggregate labour in final-good production. The calibration of final-good production then ensures that the total labour share in value added is consistent with national accounts, while the CES aggregator for energy production correctly reflects the contribution of energy capital.

et al., 2021; Hochmuth et al., 2026). Our energy production function can be rewritten in the tradition of DTC models as

$$E(t) = \left[(A_d(t)(K_d(t) + \xi))^{\epsilon(t)} + (A_c(t)K_c(t))^{\epsilon(t)} \right]^{1/\epsilon(t)}, \tag{2'}$$

where A_d and A_c are the productivities of dirty and clean capital, respectively. Thus, our specification corresponds to $A_c(t) = \varphi_c^{1/\epsilon(t)}$, $A_d(t) = \varphi_d^{1/\epsilon(t)}$ and $A_c(0)^{\epsilon(0)} + A_d(0)^{\epsilon(0)} = 1$. As a result, technical change increases productivity with decreasing marginal effects and a larger effect on clean than dirty capital.[3]

Investment adjustment costs

Investment in each capital stock is subject to convex adjustment costs. Let I_j denote investment in capital type $j \in \{f, c, d\}$. Rapid expansion of capital is costly because it requires reallocating resources and installing new capacity faster than the sector can efficiently absorb. Accordingly, the effective addition to capital is reduced by an adjustment-cost term,

$$\iota_j(t) = \chi \frac{I_j(t)^2}{K_j(t)}, \tag{3}$$

where $\chi > 0$ governs the magnitude of adjustment frictions. Adjustment costs depend on investment intensity I_j/K_j, so expanding a small sector is more costly than expanding a large one. This is important for the low-carbon transition, since clean-energy capital starts from a much smaller base than dirty capital. We will assume χ is the same in all sectors, so differences in adjustment costs depend only on the size of the sector.

Carbon intensity and stranding of dirty capital

Dirty energy capital is heterogeneous in its carbon intensity. For example, coal-fired generators emit far more per unit of capacity than combined-cycle gas turbines, and there is substantial heterogeneity in emissions intensity within fossil-fuel technologies due to differences in design, operating efficiency, and maintenance.

Let $\psi(i)$ denote the emissions intensity of type i, with $\psi'(i) < 0$, thus dirty capital types are ordered from highest to lowest emissions intensity. Let $K_d(i)$ denote the quantity of dirty capital of type i, and define $\tilde{K}_d(i) \equiv \int_0^i K_d(u)\,\mathrm{d}u$ as the cumulative quantity of dirty capital embodied in types no cleaner than i. We specify a *stranding function*, according to which emissions intensity declines exponentially with the cumulative quantity of dirty capital

[3] At equal use of clean and dirty capital ($K_d = K_c$), technical change does not affect energy unless we set $\varphi_c + \varphi_d < 1$.

embodied in successively cleaner types:

$$\psi(\tilde{K}_d(i)) = \psi_1 \exp(-\psi_2 \tilde{K}_d(i)) + \psi_3, \tag{4}$$

where the parameters (ψ_1, ψ_2, ψ_3) are calibrated to firm-level data as described in Section 4.

A key feature of the model is that we allow *negative* dirty investment, $I_d(t) < 0$, corresponding to the early retirement of existing dirty assets. When dirty capital is retired, the most carbon-intensive units (lowest i) are withdrawn first. To track early retirement, the model keeps a separate stock of stranded dirty capital, $K_s(t)$, which records the cumulative quantity of dirty capital removed from productive use.

All dirty-capital types depreciate at the same physical rate, and any new dirty investment is assumed to be of the average type of the existing distribution. In addition, we assume that emissions intensity declines exogenously at rate g_ψ, capturing gradual improvements in the carbon efficiency of dirty energy use that occur independently of climate policy. These include within-technology improvements such as higher thermal efficiency, reduced flaring and leakage, and incremental retrofits that lower emissions per unit of energy produced. Combining these elements, total emissions at time t are

$$P(t) = \exp(-g_\psi t) \int_{K_s(t)}^{K_d(t)+K_s(t)} \psi\left(\frac{K_d(0)}{K_d(t)+K_s(t)}\,\tilde{K}\right) \mathrm{d}\tilde{K}, \tag{5}$$

where the lower integration limit reflects that the dirtiest units have been stranded first, while the upper integration limit is unaffected by stranding. The average emissions intensity of the remaining productive dirty capital is therefore $\bar{\psi}(t) = P(t)/K_d(t)$.

Premature retirement of dirty capital entails an economic cost because long-lived assets are scrapped before the end of their useful lifetime. First, we conservatively assume that negative dirty investment does not release resources for consumption or other investment, i.e., dirty assets have no salvage value.[4] Second, we introduce a convex stranding-cost term to capture institutional, regulatory, and political frictions (e.g., compensation claims, regulatory delays, remediation obligations and worker displacement costs) that make large-scale or rapid

[4]When optimal dirty investment is negative, we can exclude the emergence of a secondary market with positive prices for dirty capital. This simplification is consistent with evidence of low capital resale prices in phases of forced liquidation or low demand, especially in sectors characterised by strong asset specificity (Lanteri, 2018; Kermani and Ma, 2023). We also assume zero scrap value: while *gross* scrap values are usually low but positive, these are often entirely absorbed by decommissioning costs (e.g. demolition costs, land restoration, overheads, etc.), leading to null or negative *net* scrap values. See for instance Jindal and Shrimali (2022) for coal plants in India and Raimi (2017) for power plants in the US.

asset retirement increasingly costly:

$$\Theta(t) = \mathbf{1}_{(I_d(t)<0)}\, \chi_s \frac{I_d(t)^2}{K_d(t)}, \tag{6}$$

where $\chi_s > 0$ is distinct from the adjustment cost for positive dirty investment. This specification ensures that stranding costs are zero when $I_d(t) \geq 0$, increase with the magnitude of negative investment, and are convex in $I_d(t)$.

Capital equations of motion

Final-good capital and clean energy capital evolve according to the standard accumulation equation with sector-specific adjustment costs. For $j = f$ and $j = c$,

$$\dot{K}_j(t) = I_j(t) \; - \; \iota_j(t) \; - \; \delta K_j(t), \tag{7}$$

where δ is the depreciation rate and ι_j captures sector-specific adjustment costs.

Dirty energy capital evolves similarly when investment is positive, but negative dirty investment corresponds to the premature retirement of existing dirty assets and therefore does not incur the same adjustment cost. We write

$$\dot{K}_d(t) = I_d(t) \; - \; \mathbf{1}_{(I_d(t)>0)}\, \iota_d(t) \; - \; \delta K_d(t). \tag{8}$$

When $I_d > 0$, the quadratic term captures the adjustment cost of expanding dirty capacity; when $I_d < 0$, this term vanishes and dirty capital contracts at rate I_d, with the associated economic cost of premature retirement accounted for separately through the stranding-cost function Θ in the resource constraint.

Negative dirty investment increases the stock of stranded dirty capital. Stranded assets are withdrawn from production but depreciate at the same rate as productive dirty capital:

$$\dot{K}_s(t) = \max\{-I_d(t), 0\} \; - \; \delta K_s(t). \tag{9}$$

Clean technical change

The model includes two mechanisms of clean technical change. First, exogenous improvements occur independently of the scale of clean deployment. These represent innovation spillovers from other sectors (e.g., general-purpose technologies) and basic R&D directly into clean technology that is independent of levels of clean technology use.[5] Second, endoge-

[5]This is a form of endogenous innovation that is explicitly modelled in some of the literature (Goulder and Mathai, 2000; Acemoglu et al., 2012), but Coppens et al. (2025) show that, if tractability is a concern,

nous improvements arise from learning-by-doing: as clean technologies are deployed more widely, firms accumulate experience, supply chains mature, and costs fall through economies of scale and process optimisation.

We combine these channels in the technical change factor

$$\mathrm{TC}(t) = \omega\, e^{-g_\sigma t} \;+\; (1-\omega)\left(\frac{M(t)}{M(0)}\right)^{-\varrho}, \tag{10}$$

where $\omega \in [0,1]$ determines the relative weight on exogenous vs. experience-based improvements, g_σ is the rate of exogenous technical change, and $\varrho > 0$ is the learning elasticity. Cumulative experience is proxied by cumulative abatement relative to time zero:

$$M(t) = M(0) + \int_0^t \Big(P_{\mathrm{BAU}} - P(u)\Big)\,\mathrm{d}u, \tag{11}$$

where P_{BAU} denotes emissions along a business-as-usual path without climate policy at time zero. Thus, M increases whenever realised emissions are reduced below the initial BAU level.

The technical change factor governs the substitutability between clean and dirty energy. As TC declines through innovation and learning-by-doing, it becomes easier to substitute clean for dirty energy:

$$\epsilon(t) = 1 - \frac{\mathrm{TC}(t)}{\sigma(0)}, \tag{12}$$

where $\sigma(0)$ is the calibrated initial level of substitutability. The corresponding elasticity of substitution is

$$\sigma(t) \equiv \frac{1}{1-\epsilon(t)} = \frac{\sigma(0)}{\mathrm{TC}(t)}. \tag{13}$$

Thus, as TC falls, the elasticity σ rises.

Damages and warming

The damage multiplier is

$$\Omega(t) = \exp\left(-\frac{\gamma}{2}T(t)^2\right), \tag{14}$$

where γ is the damage function coefficient and T is global mean temperature relative to its pre-industrial level.

Climate science has shown that T is approximately linearly proportional to cumulative

then it is possible to treat such R&D-driven clean technical change as exogenous provided R&D investment costs are small relative to output.

carbon dioxide emissions (IPCC, 2021), so we write

$$T(t) = \zeta S(t) = \zeta \int_0^t P(u)du, \tag{15}$$

where S stands for cumulative emissions of all greenhouse gases (GHGs) and ζ is known as the transient climate response to cumulative carbon emissions or TCRE. Although temperature responds to a CO_2 emission impulse with a delay of a few years, the relationship between temperature and cumulative emissions is still well approximated by Eq. (15) (Dietz and Venmans, 2019; Dietz et al., 2021).

Shocks

Uncertainty enters the model through four stochastic processes: shocks to the accumulation of the three types of capital, which capture economic volatility and uncertain emissions abatement costs, and temperature shocks, which capture climate uncertainty. We represent economic volatility with capital rather than TFP shocks in order to avoid adding extra state variables. All processes are assumed independent.

Each capital stock K_j follows a geometric Brownian motion:

$$\mathrm{d}K_j = (I_j - \iota_j - \delta_j K_j)\,\mathrm{d}t + \sigma_{K_j} K_j \mathrm{d}W^j, \tag{16}$$

where $\mathrm{d}W^j$ is a standard Wiener process with variance $\mathrm{d}t$ and σ_{K_j} denotes the volatility of capital type j.

Temperature is also subject to a geometric Brownian motion,

$$\mathrm{d}T = \zeta P\,\mathrm{d}t + \sigma_T T\,\mathrm{d}W^T, \tag{17}$$

where σ_T captures volatility in the climate system's temperature response to emissions. The Brownian motion $\mathrm{d}W^T$ represents both epistemic uncertainty – arising from imperfect knowledge of the climate response – and aleatoric uncertainty – random variability in the Earth system's transient behaviour.[6] Equation (17) replaces (15) in the stochastic model.

[6]This decomposition can be made explicit by assuming that epistemic uncertainty (Bayesian updating) in the TCRE parameter ζ follows the process $\mathrm{d}\zeta = \sigma_1 \mathrm{d}W_1$ and that aleatoric (weather) shocks are proportional to temperature, adding an uncorrelated geometric Brownian motion with volatility σ_2. This yields

$$\mathrm{d}T = \zeta P\,\mathrm{d}t + \frac{\sigma_1}{\zeta} T\,\mathrm{d}W_1 + \sigma_2 T\,\mathrm{d}W_2,$$

which is equivalent to the simplified expression above with $\sigma_T = \sqrt{(\sigma_1/\zeta)^2 + \sigma_2^2}$.

Consumption and utility

Aggregate consumption C is determined by the resource constraint

$$C(t) = Y(t) - I_f(t) - \mathbf{1}_{(I_d(t)>0)} I_d(t) - I_c(t) - \Theta(t). \tag{18}$$

When dirty investment is negative, existing capital is scrapped and yields no salvage value, so $-I_d(t)$ does not release resources for consumption or other investment.

To disentangle intertemporal substitution from risk aversion, we assume recursive preferences in the sense of Duffie and Epstein (1992). Let $V(x)$ denote the value function, where

$$x = (K_f, K_c, K_d, K_s, S, t) \tag{19}$$

collects the endogenous state variables, where we replace T by S since they scale linearly, and time t is an additional state argument.

Lifetime utility is given by

$$V(x) = \max_{I_f, I_c, I_d} \mathbb{E}_t \int_t^\infty f\Big(x(u), V(x(u)), I_f(u), I_c(u), I_d(u)\Big)\, \mathrm{d}u. \tag{20}$$

The instantaneous aggregator is

$$f(x, V, I_f, I_c, I_d) = \frac{[C(t)/L(t)]^{1-\eta}}{1-\eta}\, \Upsilon \;-\; (\rho - g_L)\, V\, \frac{1-\mathrm{RRA}}{1-\eta}, \tag{21}$$

with

$$\Upsilon = (\rho - g_L)\, [\,(1-\mathrm{RRA})\, V\,]^{\frac{\eta-\mathrm{RRA}}{1-\mathrm{RRA}}}. \tag{22}$$

Here η is the elasticity of marginal utility of consumption, RRA is the coefficient of relative risk aversion, and ρ is the pure rate of time preference. The pure rate of time preference is adjusted for population growth to give classical utilitarian preferences. Setting $\eta = \mathrm{RRA}$ collapses the formulation to standard expected utility, and in the deterministic case RRA is irrelevant in any case. However, in the stochastic setting the restriction that $\eta = \mathrm{RRA}$ is difficult to reconcile with empirical evidence on risky and risk-free asset returns, while in the climate-policy context it mechanically links higher risk aversion to a higher effective discount rate and therefore to weaker climate ambition.

3 Optimal capital accumulation and carbon pricing in theory

In this section, we build intuition for our quantitative results by looking at how transition frictions affect the optimal pricing/taxation of GHG emissions in theory. The results are based on solving the model as a social planning problem (Appendix A), for which we switch to effective labour units so that the model's balanced growth path is a steady state. In Appendix B, we show the decentralised competitive equilibrium equivalent to the social planner's solution.

Let λ_S denote the co-state for cumulative emissions. The co-state equation from the planner's problem is

$$\dot{\lambda}_S = (r-g)\lambda_S - c^{-\eta}\left[-y\gamma\zeta^2 S - y_e e_\epsilon \epsilon_M \frac{1}{M_0}\right], \tag{23}$$

where lowercase variables are per unit of effective labour (e.g., e for energy) and r is the consumption discount rate. Integrating this forward and defining

$$\lambda_S^{\text{SCC}} \equiv -\int_t^\infty \exp\left(-(r-g)(u-t)\right) c^{-\eta}\, y\gamma\zeta^2 S\, du, \tag{24}$$

$$\lambda_S^{\text{TC}} \equiv -\int_t^\infty \exp\left(-(r-g)(u-t)\right) c^{-\eta}\, y_e e_\epsilon \epsilon_M \frac{1}{M(0)}\, du, \tag{25}$$

we get

$$\lambda_S = \lambda_S^{\text{SCC}} + \lambda_S^{\text{TC}}.$$

Thus, the shadow price of carbon is the sum of marginal climate damages – the SCC – and the marginal social benefit of technological learning. Learning-by-doing is assumed to be an externality. Therefore, it increases the shadow price of carbon because the prospect of reducing future abatement costs creates an added incentive to abate emissions today. The additivity of λ_S^{SCC} and λ_S^{TC} comes from the fact that both depend on emissions – in the case of learning-by-doing, cumulative emissions abatement serves as a proxy in the model for experience with clean technology.[7] In a decentralised market equilibrium, the learning-by-doing externality would most plausibly be internalised by an abatement subsidy, distinct from a Pigouvian tax on emissions to internalise the climate externality λ_S^{SCC} (Appendix B).

[7] Coppens et al. (2025) compare this case with others where the two externalities are not additive.

MAC decomposition with adjustment costs

Define the investment-capital ratio $a_j \equiv i_j/k_j$, the investment growth rate $g_j^I \equiv \dot{i}_j/i_j$, and the capital growth rate $g_j^K \equiv \dot{k}_j/k_j$ for capital type j. Let y_{k_d} and y_{k_c} be the marginal products of dirty and clean capital, respectively, and P_{k_d} be the marginal emissions from dirty capital.

The marginal abatement cost is defined in consumption units as $\text{MAC} \equiv -\lambda_S c^\eta$. The specific form of the MAC depends on whether the economy is in a stranding regime, in which dirty capital is being retired prematurely, or not. Starting with the latter case $i_d \geq 0$, combining the Euler equations for clean and dirty capital yields

$$\begin{aligned} \text{MAC} \cdot (1 - 2\chi a_d)\ P_{k_d} = & \underbrace{(1 - 2\chi a_d)\, y_{k_d} - (1 - 2\chi a_c)\, y_{k_c}}_{\text{(A) Net marginal product gap}} \\ & + \underbrace{\chi(a_d^2 - a_c^2)}_{\text{(B) Current adjustment-cost pressure}} \\ & + \underbrace{\frac{g_d^I - g_d^K}{\frac{k_d}{2\chi i_d} - 1} - \frac{g_c^I - g_c^K}{\frac{k_c}{2\chi i_c} - 1}}_{\text{(C) Dynamic adjustment effects}} . \end{aligned} \tag{26}$$

The right-hand side of (26) is the excess return on dirty capital relative to clean at the social optimum, including marginal products, adjustment costs, and dynamic considerations. Dividing this by the term $(1 - 2\chi a_d)P_{k_d}$, which reflects both the marginal emissions from dirty capital and the wedge introduced by adjustment costs, yields the carbon price required to make this return differential consistent with optimal investment allocation.

(A) Net marginal product gap. The term

$$(1 - 2\chi a_d)y_{k_d} - (1 - 2\chi a_c)y_{k_c}$$

is the difference in the effective marginal products of dirty and clean capital. It measures the direct loss from shifting one unit of capital from dirty to clean. Without adjustment costs, the MAC is simply $y_{k_d} - y_{k_c}$. With adjustment costs, since clean investment intensity a_c is high relative to a_d along the low-carbon transition, the direct effect of adjustment costs is to increase the MAC.

(B) Current adjustment-cost pressure. The term

$$\chi(a_d^2 - a_c^2)$$

arises because adjustment costs scale with the square of investment intensity. Holding investment fixed, increasing the capital stock reduces adjustment costs by increasing the denominator. Since clean investment intensity a_c is high along the transition, an increase in clean capital reduces adjustment costs more strongly than an equivalent increase in dirty capital. As a result, the clean-side term χa_c^2 dominates, so $\chi(a_d^2 - a_c^2) < 0$. Intuitively, expanding clean capital relaxes adjustment-cost pressures more than expanding dirty capital, making the reallocation toward clean activities easier at the margin and thereby lowering the MAC.

(C) Dynamic adjustment effects. The term

$$\underbrace{\left(g_d^I - g_d^K\right) \frac{2\chi a_d}{1 - 2\chi a_d}}_{\text{dirty-side dynamic effect}} - \underbrace{\left(g_c^I - g_c^K\right) \frac{2\chi a_c}{1 - 2\chi a_c}}_{\text{clean-side dynamic effect}} \tag{27}$$

captures how the time profile of adjustment costs affects the current return differential between dirty and clean capital. These terms arise because adjustment costs are not only a function of current investment intensities, but also of how those intensities are changing over time. If dirty investment is falling relative to the dirty capital stock, so that $g_d^I - g_d^K < 0$, then future dirty adjustment costs are declining. This makes current dirty investment less attractive relative to the future, which increases the MAC. If clean investment is growing faster than the clean capital stock, so that $g_c^I - g_c^K > 0$, then part of the required clean expansion is already being brought forward. This helps smooth future adjustment costs and therefore lowers the MAC required to tilt investment toward clean capital.

MAC under a stranding regime and the role of stranding costs

When $i_d < 0$, the planner strands high-emissions capital. Let λ_{k_s} be the co-state on stranded capital. The marginal resource cost of stranding is

$$\lambda_{k_d} - \lambda_{k_s} = 2\chi_s \, a_d \, c^{-\eta}. \tag{28}$$

Using the dirty-capital Euler equation together with (28), we obtain

$$\begin{aligned}\text{MAC} \cdot (P_{k_d} - P_{k_s}) = & \underbrace{y_{k_d}}_{\text{(A}') \text{ Marginal product of dirty capital}} \\ & + \underbrace{\chi_s a_d^2}_{\text{(B}') \text{ Level effect of stranding costs}} \\ & + \underbrace{-2\chi_s a_d \left(r + g_L + \eta \frac{\dot{c}}{c} + \delta - g_d^I + g_d^K \right)}_{\text{(C}') \text{ Dynamic stranding effects}}.\end{aligned} \tag{29}$$

The term $P_{k_d} - P_{k_s}$ is the marginal emissions from keeping one unit of dirty capital in operation rather than stranding it, expressed in effective emissions units consistent with the MAC definition.

(A′) Net marginal product gap under stranding. When $i_d < 0$, the relevant marginal cost of increasing the rate of stranding is just the foregone output from dirty capital, given by its marginal product.

(B′) Stranding costs and current strain. In the stranding regime the quadratic term in the MAC is

$$\chi_s a_d^2,$$

which reflects the convex cost of disinvesting from dirty capital. Because stranding implies negative investment, a large negative a_d means that disinvestment is occurring at high intensity relative to the existing dirty capital stock. With quadratic costs, this makes further increases in the rate of stranding disproportionately costly, which raises the MAC.

(C′) Dynamic stranding effects. The term

$$-2\chi_s a_d \left(r + g_L + \eta \frac{\dot{c}}{c} + \delta - g_d^I + g_d^K \right)$$

is the dynamic counterpart to (C). It captures how the time profile of stranding costs affects the current marginal cost of stranding. Because $a_d < 0$, this term is positive when dirty investment is falling relative to the dirty capital stock, so that $g_d^I - g_d^K < 0$. In that case the local dynamics imply a steeper path of stranding intensity, which raises the marginal cost of additional stranding today and therefore raises the MAC. Intuitively, when the optimal path requires rapid stranding, the planner has an incentive to smooth those costs over time rather than concentrate them immediately.

Stochastic extension: risk-adjusted carbon pricing

With recursive preferences and stochastic shocks, the optimal carbon price is obtained by applying the envelope theorem to the HJB equation and converting the shadow value of cumulative emissions from utility units into consumption units. Denoting by V_S the derivative of the value function with respect to cumulative emissions (equal to λ_S in the above analysis), and by V_{k_j} the derivatives with respect to the capital stocks (λ_{k_j} above), the optimal carbon price can be written as

$$\begin{aligned}
\mathrm{MAC} = \Phi\Bigg\{ & \underbrace{\Psi c^{-\eta} y \gamma \zeta^2 S - V_{SS} P}_{\text{(1) Marginal damage cost}} \\
& + \underbrace{\Psi c^{-\eta}\left[y_e e_\epsilon \epsilon_M \frac{1}{M(0)}\right] - V_{tS}}_{\text{(2) Technological-change effect}} \\
& \underbrace{\begin{aligned}
& - V_{Sk_f}\left[i_f - \chi\frac{i_f^2}{k_f} - (\delta + g_L + g)k_f\right] - V_{Sk_d}\left[i_d - \mathbf{1}_{i_d>0}\chi\frac{i_d^2}{k_d} - (\delta + g_L + g)k_d\right] \\
& - V_{Sk_c}\left[i_c - \chi\frac{i_c^2}{k_c} - (\delta + g_L + g)k_c\right] - V_{Sk_s}\left[-\mathbf{1}_{(i_d<0)} i_d - (\delta + g_L + g)k_s\right]
\end{aligned}}_{\text{(3) Capital reallocation effects}} \\
& \underbrace{\begin{aligned}
& - V_{SS} S \sigma_T^2 - \frac{1}{2} V_{SSS} S^2 \sigma_T^2 \\
& - \frac{1}{2} V_{Sk_d k_d} k_d^2 \sigma_{K_d}^2 - \frac{1}{2} V_{Sk_c k_c} k_c^2 \sigma_{K_c}^2 - \frac{1}{2} V_{Sk_f k_f} k_f^2 \sigma_{K_f}^2
\end{aligned}}_{\text{(4) Risk premia}} \Bigg\}.
\end{aligned} \tag{30}$$

where

$$\Phi = \frac{(1-\eta)}{\hat{\rho}\left[1 - \mathrm{RRA} + (\mathrm{RRA} - \eta) c^{1-\eta}\left((1-\mathrm{RRA})V\right)^{-\frac{1-\eta}{1-\mathrm{RRA}}}\right]} \frac{1}{V_{k_c}\left(1 - 2\chi\frac{i_c}{k_c}\right)}, \tag{31}$$

$$\Psi(V) = \frac{\hat{\rho}}{\left((1-\mathrm{RRA})V\right)^{\frac{1-\eta}{1-\mathrm{RRA}}-1}}. \tag{32}$$

Here $\hat{\rho}$ is the effective discount rate in the recursive aggregator.

To interpret the stochastic carbon price in Eq. (30), notice that $-V_S$ is the optimal carbon price in utility units and is positive. This sign convention is useful in reading the various cross-derivatives. Whenever a marginal emission raises future damages or increases marginal utility in adverse states, it contributes positively to the optimal carbon price. The individual components of Eq. (30) can then be read as follows.

(1) Marginal damage cost. The first term generalises the deterministic social cost of carbon. The component $\Psi c^{-\eta} y \gamma \zeta^2 S$ captures the current marginal damage of an additional emission through higher temperatures and lower output, scaled by recursive preferences. The term $-V_{SS}P$ reflects how a marginal emission changes the future path of marginal damages. On a path of rising temperatures, a convex damage function implies increasing marginal damages, so that $-V_{SS} > 0$.

(2) Technical change effects. The second set of terms captures how emissions affect future abatement costs through technical change. The first element reflects the marginal social benefit of learning, just as in the deterministic model. The second element, V_{tS}, embodies the exogenous forces in the model that reduce the MAC, specifically the time-dependence of the elasticity of substitution between clean and dirty capital in Eq. (10) and the time-dependent emissions intensity of dirty capital. These effects lower the optimal carbon price by making abatement exogenously cheaper. In the deterministic model, such trends are included in $\dot{\lambda}_S$. In the stochastic model, the value function depends directly on time, and the envelope condition for cumulative emissions yields the additional V_{tS} term.

(3) Capital reallocation effects. These terms capture how emissions affect the opportunity cost of holding different types of capital and hence the optimal reallocation between dirty and clean activities. The derivatives V_{Sk_j} measure how a marginal emission changes the value of each capital stock through its impact on future damages and marginal utility. The drift terms in square brackets describe the expected evolution of final-good, dirty, clean and stranded capital net of depreciation and, where relevant, adjustment costs. In the presence of adjustment costs, these terms closely parallel the deterministic MAC decomposition shown above.

(4) Risk premia. The main change to the optimal carbon price in a stochastic setting is the introduction of a set of risk premia. The first two terms $-V_{SS}S\sigma_T^2 - \frac{1}{2}V_{SSS}S^2\sigma_T^2$ constitute a temperature risk premium. Uncertainty in the temperature response to emissions implies that a marginal tonne of CO_2 increases not only expected damages but also the dispersion of future temperature outcomes. With risk aversion, greater downside climate risk raises the value of mitigation because marginal damages are weighted more highly in low-consumption states. The next three terms constitute capital risk premia. Dirty, clean and final-good capital are risky assets whose returns covary with marginal utility, changing the distribution of future consumption.

It is useful to contrast this expression with a straw-man model in which there are no adjustment costs, no endogenous technical change, and no uncertainty, but in which emissions arise from capital use in the standard way. In such a model, the optimal carbon price

reflects expected marginal damages and the frictionless opportunity cost of emissions through their effect on future output and capital values. Relative to this benchmark, adjustment costs change the timing and marginal cost of capital reallocation by making rapid shifts between dirty and clean capital costly, endogenous technical change links emissions to future abatement costs, and uncertainty adds a set of risk premia.

4 Calibration

We calibrate a discrete-time version of the model with annual time steps, with numerical simulations starting in 2025. Capital stocks and output are measured in trillion 2024 US dollars. We begin with the most novel elements of our calibration: (i) the stranding function; and (ii) adjustment and stranding costs. We then present the remaining parameters and initial conditions, drawn from standard values in related literature or calibrated to match recent empirical evidence and key features of the initial state. Table A1 in Appendix C.1 summarises all parameter values.

Emissions intensity and stranded assets

Dirty capital K_d comprises a continuum of heterogeneous polluting technologies i with emissions intensity $\psi(i)$, as in Equation (4). The key objective of this subsection is to discipline the model's stranding function $\psi(\tilde{K}_d(i), i)$, i.e., the relationship governing which dirty assets are selected for early retirement and how much emissions abatement this achieves. To do so, we conduct a novel analysis of company-level data on tangible assets and emissions, complemented by plant-level evidence that confirms our benchmark patterns. This allows us to recover a realistic distribution of capital emissions intensities, taking into account that assets differ not only in their current emissions intensity, but also in the share of emissions that remains embodied in their residual lifetime. Appendix C.2 provides further details supporting what follows.

We start by considering all active companies operating in a set of particularly polluting sectors covered by Bureau van Dijk's Orbis database. We include 22 productive NACE sectors. From Orbis, we extract balance-sheet information on the monetary value of physical assets potentially at risk of stranding. We treat all types of tangible assets except land as strandable (buildings, plant and machinery, transport equipment, leased assets and other plant, property and equipment). We then match these data to company-level Scope 1 GHG emissions from the LSEG financial database. This yields a sample of 1,983 companies, emitting roughly 9.3 $\mathrm{GtCO_2e}$. To account for incomplete coverage of sectoral emissions, we compute sector-specific scaling factors using data from Crippa et al. (2025) and UNFCCC

(2026), and apply them to both emissions and tangible-asset values. The scaled sample accounts for 33.2 $GtCO_2e$, or about 64.2% of global emissions. We attribute the remaining emissions to (i) other productive sectors; and (ii) end-use household activities.

The age and expected lifetime of capital assets are important. When productive capital is stranded, society avoids not only current emissions, but also those that would have been released over the asset's remaining life (its 'embodied' emissions). To incorporate this temporal dimension, we exploit the availability on Orbis of both 'net' and 'gross' physical asset values, denoted by K_n and K_g respectively. Assuming linear depreciation, as is common in corporate accounting (Jackson et al., 2009), we proxy asset age as $a = (1 - K_n/K_g)\, LT$, where LT is the typical lifetime of tangible assets. We then fix $LT = 34.65$ years jointly with our depreciation rate $\delta = 0.04$.

We define 'forward-looking emissions intensity' (FLEI) as the annual emissions intensity of a company, rescaled by a factor that reflects the remaining lifetime of the asset stock. This adjustment is needed to reconcile the exponential depreciation structure of the model with the linear depreciation embedded in the Orbis company data. Equalising embodied emissions across the two depreciation structures, i.e., setting $P_{Orbis}(LT - a) = P_{model}/\delta$, yields

$$\psi(i) = \frac{P(i)}{K_d(i)} \frac{LT - a(i)}{1/\delta}. \tag{33}$$

After dropping observations with emissions below 1 $ktCO_2e$ and winsorising the FLEI distribution at the 2.5th and 97.5th percentiles to mitigate the influence of outliers, we obtain a sample of 1,849 companies with observed emissions and tangible assets, scaled by sector-specific factors.

We rank companies by their FLEI and obtain the upper panel of Figure 1. Each bar represents a company, with width proportional to the value of its tangible assets. Assets with high FLEI – on the left of the plot – are stranded first when $i_d < 0$, either because the underlying technology is particularly emissions-intensive or because a large fraction of emissions remains locked into the asset's residual life. It is from this calibration that we are able to identify a 'dirty tail' of assets in which emissions are highly concentrated: the first \$10 tn of dirty capital accounts for almost 70% of the emissions in our company sample.

To see the implications of this concentration of emissions for abatement costs, define 'stranding cost per cumulative emission' (SCCE) as the cost of abating one ton of CO_2e by decommissioning company i's dirty assets, taking into account both current and future emissions:

$$\text{SCCE}(i) = \frac{K_n(i)}{P(i)(LT - a(i))}. \tag{34}$$

The lower panel of Figure 1 ranks companies by their SCCE, with bar width proportional to

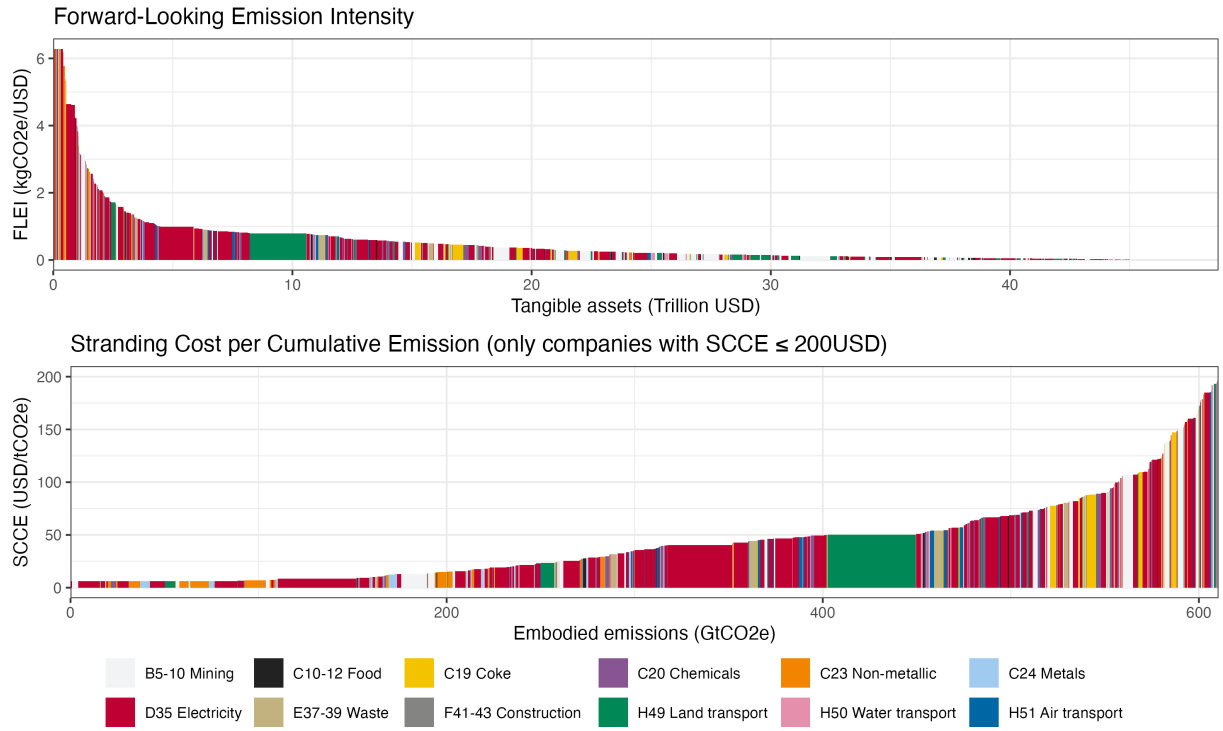

Figure 1: Company-level forward-looking emissions intensity (FLEI - upper panel) and stranding cost per cumulative emission (SCCE - lower panel). Total observations: 1,849. Sector-specific scaling factors are applied. The width of each bar is proportional to company physical assets (upper panel) or the amount of embodied emissions (lower panel). FLEI/SCCE values are winsorised at 95%. Companies with emissions lower than 1 ktCO_2e are dropped. Only companies with SCCE $\leq$ \$200 are shown in the lower panel.

their scaled embodied emissions, i.e. the emissions their dirty tangible assets would generate if operated until the end of their natural lifetime. For readability, we display only companies with an SCCE below \$200. Assets with low SCCE values – on the left of the plot – are the cheapest to strand. These are the assets with a high FLEI. A substantial amount of emissions can be abated at limited monetary loss, for example, more than 400 GtCO_2e at a cost below \$50/tCO_2e.

After having estimated companies' FLEI, we fit a shifted exponential function to the empirical curve, $\psi(\tilde{K}_d(i)) = \psi_1 \exp(-\psi_2 \tilde{K}_d(i)) + \psi_3$, where ψ_3 is the minimum emissions intensity, $\psi_1 + \psi_3$ is the intercept at $\tilde{K}_d = 0$, and ψ_2 governs the rate at which intensity declines as $\tilde{K}_d$ increases.[8] We obtain $\psi_1 = 6.120$, $\psi_2 = 0.729$, and $\psi_3 = 0.767$, which we use to calibrate the stranding cost function in Equation (4).

Given the importance of the FLEI curve and the dirty tail, we test the robustness of our

[8]To improve the fit in the initial part of the FLEI curve, where the model selects firms to strand, we apply a logistic weighting pattern and estimate the function using DEoptim (Mullen et al., 2011).

calibration by repeating the analysis with global coal and gas plant data from GEM (2026); results are reported in Appendix C.3. Despite the different data source and methodology – plant monetary values are approximated using region-specific overnight capital cost estimates rather than balance-sheet data – the plant-level FLEI and SCCE distributions closely replicate the company-level patterns. Coal and gas plants occupy the far left of both curves in Figure 1, confirming that they represent the most emissions-intensive and cheapest-to-strand assets. Consistent with the company-level evidence, substantial near-term abatement can be achieved through the targeted retirement of a comparatively small set of very dirty assets: almost the entire global coal and gas power fleet – representing close to 250 $\mathrm{GtCO_2e}$ – can be stranded at an SCCE below $\$20/\mathrm{tCO_2e}$.

Finally, besides endogenous abatement through capital stranding, we allow emissions intensity to decline exogenously at rate g_ψ. We set $g_\psi = 0.01/\mathrm{yr}$ using World Bank (2026) data on total GHG emissions (net of agricultural methane and nitrous oxide emissions) and value added in industry, including construction. The rate is computed over 1994–2007 to limit the influence of climate policies on the underlying trend.

Adjustment costs

Our model has two key parameters limiting the speed of the energy transition. The first, χ, captures constraints on how quickly productive capital of any sort can be scaled up. These constraints will be most relevant to the clean-energy sector, since it starts small but must expand rapidly. The second, χ_s, captures frictions impeding the premature retirement of dirty capital.

Since χ is most relevant to the clean-energy sector, our calibration strategy anchors it to historical episodes of unusually rapid expansion in large-scale, capital-intensive energy systems. These episodes provide a benchmark for how quickly new capacity can be deployed under real-world constraints. We map these peak net growth rates of the energy capital stock into χ. From Equations (3) and (7), the peak net growth rate of capital can be written as

$$g_{\text{peak}} = a - \chi a^2 - \delta,$$

where $a \equiv i/k$ as before. The sector subscript is suppressed given that we set χ equal across sectors. Let the marginal effectiveness of clean investment be

$$\text{eff} \equiv 1 - 2\chi a,$$

which measures the fraction of an additional unit of investment that is transformed into productive capital. Rather than interpreting observed peak expansion episodes as operating

at zero marginal effectiveness, we impose a conservative effectiveness floor $\text{eff} \geq \text{eff}_{\min}$. This reflects the idea that even during intense mobilisation, a substantial fraction of marginal investment continued to translate into productive capacity.

Imposing both the peak-growth and effectiveness-floor conditions yields the closed-form expression

$$\chi = \frac{1 - \text{eff}_{\min}^2}{4\Big(g_{\text{peak}} + \delta\Big)}.$$

Assuming $\text{eff}_{\min} \in [0.5, 0.7]$ and $\delta = 0.04$, we estimate χ based on empirical data on g_k^{peak}, which we obtain from three historical episodes of rapid energy-capital build-out: the global nuclear expansion in the mid-1980s (Bennett and Skjoeldebrand, 1986), China's power-system expansion in the late 2000s (Nan and Moseley, 2011), and the US combined-cycle gas turbine boom of the early 2000s (EIA, 2025b). These data indicate $g_{\text{peak}} \in [0.10, 0.12]$, implying $\chi \in [0.8, 1.35]$. We therefore set $\chi = 1$ as our baseline value.[9]

We calibrate χ_s through a similar but separate procedure, which uses historical evidence from rapid phase-outs of coal-fired electricity. These transitions provide direct information about the upper-tail rates at which modern power systems have been able to decommission dirty capital under real-world constraints. Let $g_{\text{dec}} > 0$ denote the net proportional rate of dirty-capital decline, so that $\dot{k}_d/k_d = -g_{\text{dec}}$. When dirty investment is negative, define the retirement intensity

$$u \equiv \frac{|i_d|}{k_d}.$$

Given the law of motion for dirty capital,

$$u = g_{\text{dec}} - \delta.$$

Evidence from coal phase-outs in the United States, United Kingdom and Germany suggests upper-tail values of $g_{\text{dec}} \in [0.05, 0.07]$. With $\delta = 0.04$, this implies $u \in [0.01, 0.03]$ (EIA, 2025b, 2023; BEIS, 2024; WSB, 2019).

Using Equation (6), an economically interpretable object is the additional frictional cost per unit of prematurely retired capital,

$$\kappa \equiv \frac{\Theta}{|i_d|} = \chi_s u.$$

The parameter κ captures real resource costs associated with early retirement, such as com-

[9]The literature uses a very wide range, from below 0.1-0.2 (e.g. van der Ploeg and Rezai, 2020a; David and Venkateswaran, 2019) to above 10 (Pindyck and Wang, 2013; Hambel et al., 2024). Our value is similar to Asker et al. (2014).

pensation schemes, remediation obligations, and labour-transition programmes, over and above the mechanical destruction of productive capital. Evidence from coal power-plant decommissioning suggests that physical retirement costs alone range from a few percent of replacement value to over 10% in some cases (Raimi, 2020; EIA, 2025a). Because κ is intended to capture a broader set of frictional costs than physical decommissioning alone, we conservatively calibrate $\kappa \in [0.05, 0.15]$. Using midpoint values $u = 0.02$ and $\kappa = 0.10$ implies a central calibration of $\chi_s = 5$.

Clean technical change

Clean technical change is governed by Equation (10), according to which it is a convex combination of exogenous, time-dependent technical change, and learning-by-doing. Our calibration is based on data from Coppens et al. (2025), who use large-scale evidence from climate-mitigation scenarios[10] to recover predicted rates of clean technical change across IAMs. Coppens et al. (2025) estimate that exogenous technical change reduces the slope of the MAC curve by 2.7% per year in their preferred specification, which we use as our central value for g_σ. For learning-by-doing, they estimate a learning elasticity of about 0.21, implying a learning rate of roughly 10%: each doubling of cumulative abatement reduces the MAC slope by 10%. Choosing $\varrho = 0.1$ results in the same learning rate for the marginal productivity of clean capital in our model and choosing $M(0) = 50\,\mathrm{GtCO_2}$ makes this learning rate stable over time.[11] Finally, the parameter ω governs the relative weight of exogenous and endogenous technical change. Since the evidence does not identify this parameter, we follow the principle of insufficient reason and set $\omega = 0.5$, assigning equal weight to the two channels. This is broadly consistent, however, with evidence that a substantial share of clean innovation is driven by spillovers and complementarities with innovation outside the focal technology (Taalbi, 2020). As a result, the elasticity of substitution σ reaches 5.02 in 2100, which is at the upper end of what is used in the literature.

Energy production

We calibrate the initial stock of dirty capital $K_d(0)$ using the same empirical analysis employed to construct the stranding cost function. The scaled company-level Orbis data imply total dirty capital of \$47.2 tn. However, this figure omits some relevant sectors such as agri-

[10] They combine the databases of the Intergovernmental Panel on Climate Change Sixth Assessment Report and the Network of Central Banks and Supervisors for Greening the Financial System (IPCC, 2022; NGFS, 2024).

[11] More particularly, the implied learning elasticity of clean productivity with respect to cumulative experience is 0.20 at initial capital levels, rising to an average of 0.23 over 2025–2050 and 0.21 over the full 2025–2100 horizon.

culture and households, and does not capture the fact that firms in 'clean' sectors employ small amounts of dirty capital in reality. Instead, we set $K_d(0) = \$58.4\,\text{tn}$, so that initial emissions $P(0)$ match the area under the fitted emissions-intensity curve $\psi(i)$ at $t = 0$. We set $P(0) = 53.2\,\text{GtCO}_2\text{e}$, equal to global GHG emissions in 2024 according to EDGAR (Crippa et al., 2025). The shifter parameter $\xi = 0.01$ is set to a small but non-zero value.

We set $K_c(0)$ by targeting the share of low-emission sources – renewables, nuclear, and fossil fuels with CCUS – in total final energy consumption, which was 16% in 2024 (IEA, 2025). The share of clean sources in electricity generation/capacity is much higher, but our energy sector represents all productive activities that use energy while generating GHG emissions, not just the power sector. Ideally, we would target the share of clean *capital* rather than *output*, but comprehensive capital-stock data are only available for the power sector. We therefore use the clean share in final energy consumption as the most appropriate empirical target and set $K_c/(K_c + K_d + \xi) = 0.16$.

We set the initial elasticity of substitution between energy types to $\sigma(0) = 2$, implying $\epsilon(0) = 0.5$. This value is in line with the capital-based estimates in Papageorgiou et al. (2017) and the rest of the related literature, where values usually range from 1.5 to 4, indicating a relatively high degree of substitutability between clean and dirty energy.

The profit-maximising condition in the energy sector, which equalises the marginal products of the two energy-capital stocks, then implies CES share parameters $\varphi_d = 0.696$ and $\varphi_c \equiv 1 - \varphi_d = 0.304$. Appendix C.4 provides more details.

Final good production

We calibrate the starting value of TFP $A(0) = 8.09$ so that, given initial values for labour and the capital stock, the 2024 value of global GDP is $\$110.98\,\text{tn}$ (World Bank, 2026). As is common in the literature, we assume TFP grows at a long-run rate $g_A = 2\%/\text{yr}$. We set the initial value of population $L(0) = 8.16\,\text{bn}$ (UN, 2024), and we set population growth equal to the average growth rate in the UN's medium variant scenario for 2020-2100, i.e. $g_L = 0.29\%/\text{yr}$. The share of capital in production α is assumed to be 0.3, as customary in the literature.

We initialise the final-good capital stock to $K_f(0) = \$321.3\,\text{tn}$ so that the economy starts on a balanced growth path (see Appendix C.4 for details). This is important because we want transition dynamics to be driven by changes in endogenous factors rather than by imbalanced initial capital stocks. In the model expressed in effective labour units, balanced growth corresponds to a steady state. We therefore calibrate the initial capital stock from a 'BAU steady state', in which emissions no longer generate additional warming and temperature remains at its initial level, so that the economy is interpreted as having already adjusted to

the damages associated with past warming.

We then set the CES share parameter for energy, φ_E, so that the marginal products of final-good and energy capital are equalised. The CES share parameter for final-good capital is then given by $\varphi_f \equiv 1 - \varphi_E$. Finally, we conservatively set the elasticity of substitution $\sigma_f = 0.2$, which implies $\epsilon_f = -4$, in line with the short-run estimates of Olovsson and Vestin (2023) and Hassler et al. (2021).

Damages and warming

WMO (2025) estimate that climate-induced warming in 2024 relative to the 1850-1900 pre-industrial baseline lies between 1.34 °C and 1.41 °C. We take the midpoint of 1.38 °C as our initial value. IPCC (2021) reports a central estimate of the TCRE $\zeta = 0.00045\,°\text{C}/\text{GtCO}_2$. However, this estimate does not take into account the zero-emissions commitment or some known climate feedbacks. After incorporating these effects, we set $\zeta = 0.0006\,°\text{C}/\text{GtCO}_2$.[12] Finally, we set the damage parameter $\gamma = 0.0201$, so that the climate damage function matches the meta-analytic estimate in Howard and Sterner (2017), including both catastrophic and non-catastrophic damages.

Shocks

We calibrate the temperature uncertainty parameter σ_T to replicate the 80% confidence interval of temperatures under the IPCC's RCP2.6 and RCP4.5 warming scenarios in 2100. This gives $\sigma_T = 0.0265\,°\text{C}/\text{yr}^{1/2}$. We choose scenarios RCP2.6 and RCP4.5 because they span the optimal emissions in our model. We impose a common capital volatility parameter $\sigma_{K_j} = 0.015/\text{yr}^{1/2}\ \forall j$. This is based on van den Bremer and van der Ploeg (2021), who calibrate capital shocks to the standard deviation of global GDP growth over 1961–2015. As they note, this is a conservative choice: asset return volatility is empirically much larger, and calibrating to it would generate a higher risk premium on the optimal carbon price. Imposing common capital shocks enables us to disentangle their effects more straightforwardly.

Preferences

We use the expert survey of Drupp et al. (2018) to set both the elasticity of marginal utility of consumption (i.e., the inverse of the elasticity of intertemporal substitution) $\eta = 1.35$ and

[12] Emissions in our model include all GHGs, converted to CO_2 using their 100-year Global Warming Potentials. Therefore, we need an estimate of the TCRE based on an analysis of the temperature response to GHG emissions excluding non-CO_2 forcers. According to IPCC (2021), the median cumulative CO_2 budget for 0.93°C warming from 2020 onwards is 1350 $GtCO_2$. This includes 0.1-0.2°C warming from non-CO_2 forcers, which should be excluded. Taking the midpoint, the 1350 $GtCO_2$ results in 0.78°C extra warming. Rounding the TCRE parameter to four decimal places, we get $\zeta = 0.0006$.

the utility discount rate $\rho = 1.1\%/\text{yr}$. The coefficient of relative risk aversion RRA is set equal to 4, following Barro (2009).

5 Deterministic results

In this section, we present results from the deterministic model, reporting the optimal paths under base parameter settings as well as a sensitivity analysis. We also explore the roles of adjustment costs and capital with very high emissions intensity. The model is solved numerically in Julia as a finite-horizon nonlinear optimisation problem over investment paths, using automatic differentiation to compute gradients. We employ a multi-start local search strategy (combining L-BFGS with bound constraints) to mitigate non-convexities, and verify robustness by comparing solutions across alternative initialisations and solver settings.

The optimal transition

Figure 2 reports optimal emissions, temperatures, the SCC, capital stocks, energy demand and growth for the deterministic model. The key result is that despite imposing substantial adjustment costs and limiting the substitutability of dirty energy in the short run, optimal emissions fall steeply (panel a). Using the half life of emissions – i.e., the time taken for emissions to fall to half of their starting level – as a measure of the speed of the transition, emissions halve in 10 years. Thereafter, the rate of emissions decrease slows, but optimal abatement limits the increase in temperature to 2.1 °C in 2100 (panel b).

The SCC – the Pigouvian carbon tax – starts at \$181/tCO_2e and rises to \$331/tCO_2e in 2050 (panel c). Together with an additional abatement subsidy to internalise the learning externality, this incentivises immediate disinvestment from the dirtiest energy capital, after which it is optimal to let dirty capital depreciate at its base rate of 4%. \$2.3 tn of dirty capital is stranded in the first year alone, and the stock of stranded capital reaches \$4 tn after seven years (panel d), equivalent to 1% of total capital and 7% of dirty capital. Referring back to Figure 1, readers will notice that this corresponds to the point at which the FLEI curve flattens out. Thus, optimal disinvestment from dirty capital quickly removes the dirty tail, then stops. Dirty capital is gradually replaced by clean capital, with the latter stock overtaking the former in 2036. In 2100, dirty capital accounts for less than 3% of energy capital.

In addition to substituting dirty capital for clean, the economy adjusts to the higher cost of energy by reducing energy demand in the short run (panel e). Indexed to 100 in 2024, optimal energy demand falls by around 10% by the early 2030s, before resuming strong upwards growth. Optimal output/consumption grow all along the transition path (panel f),

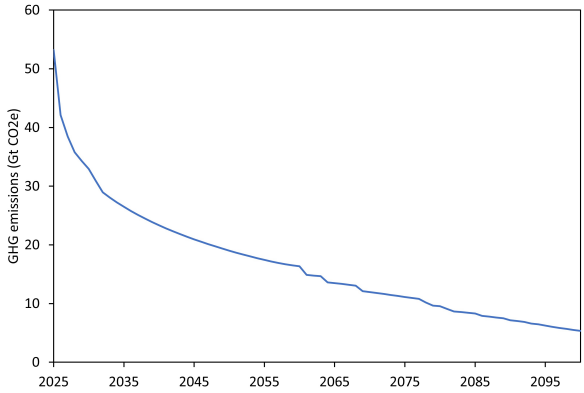

(a) Emissions P

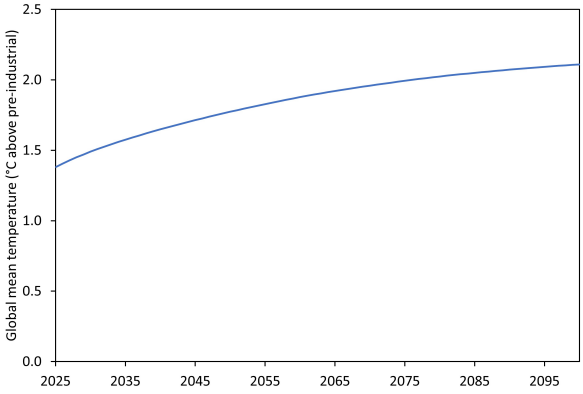

(b) Temperature T

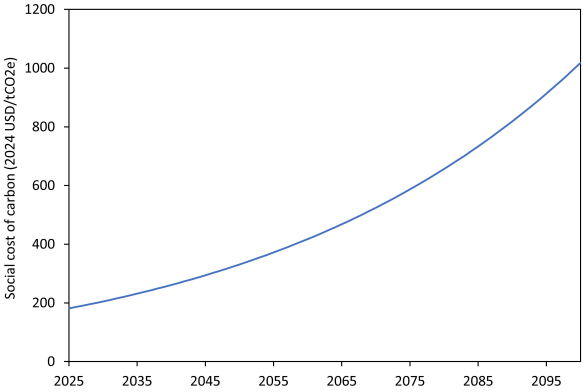

(c) SCC/Pigouvian carbon tax

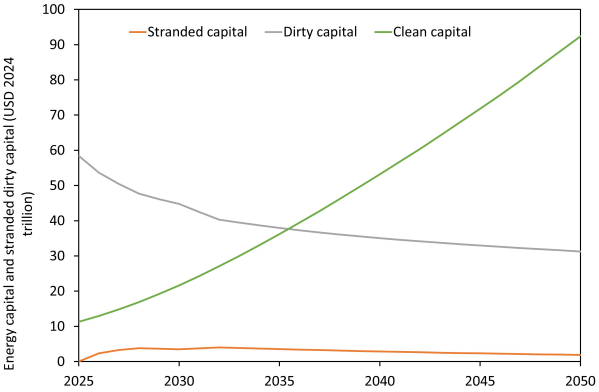

(d) Energy/stranded capital, K_c, K_d and K_s

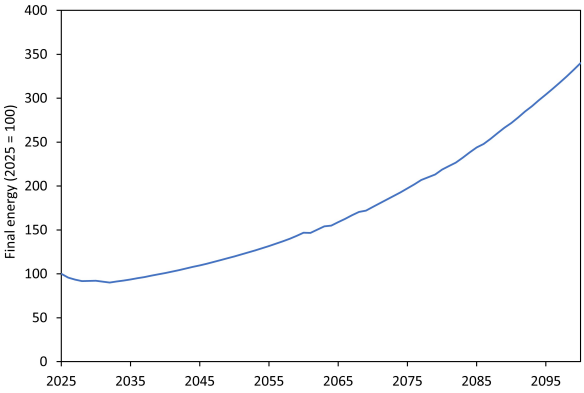

(e) Energy E

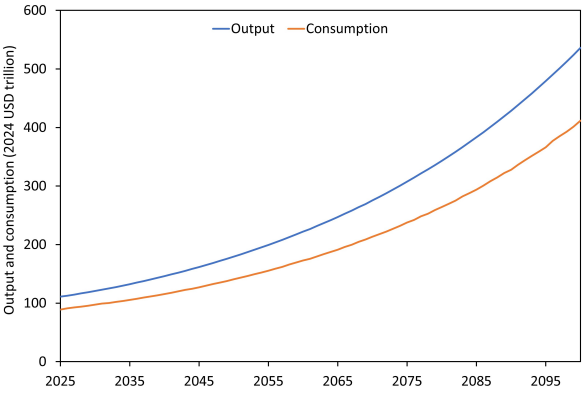

(f) Output Y and consumption C

Figure 2: Optimal transition paths.

with an initial growth rate of around 1.7%, increasing to around 2.2% in the second half of the century.

Sensitivity analysis

Table 1 reports the results of a one-factor-at-a-time sensitivity analysis. For each factor, we report two key outputs of the model, which we interpret as sufficient statistics for (i) the speed of the initial transition and (ii) the long-run ambition of climate policy. For (i) we use the emissions half-life. For (ii) we use the increase in global mean temperature in 2100. At default parameter settings these are 10 years and 2.1 °C, respectively.

Three broad types of relationship emerge. First, there are parameter variations that primarily affect the speed of the initial transition, but not the long-run ambition of climate policy. These include variations in adjustment costs χ and χ_s, the capital depreciation rate δ, the elasticity of substitution between final-good capital and energy σ_f, and the TCRE ζ. Increasing adjustment costs slows the reallocation of capital and lengthens the emissions half-life, while removing adjustment costs accelerates the transition. Varying the depreciation rate mainly affects the optimal transition through its effect on dirty capital: lower depreciation increases the persistence of dirty capital and slows the transition, while higher depreciation speeds it up through faster capital turnover. Reducing the elasticity of substitution between final-good capital and energy makes it more costly to reduce energy use in production and lengthens the emissions half-life, whereas higher substitutability has the opposite effect. A low TCRE reduces the climate response to emissions, lowering marginal climate damages and weakening incentives for abatement. This increases the emissions half-life. Conversely, a high TCRE raises marginal damages and strengthens incentives for abatement, shortening the half-life. However, temperature in 2100 changes much less than emissions do, because the direct effect of TCRE on warming is partly offset by the endogenous response of optimal cumulative emissions: a higher TCRE induces stronger abatement and lower cumulative emissions, while a lower TCRE induces weaker abatement and higher cumulative emissions.

The second type of relationship encompasses parameter variations that affect both the long-run ambition of climate policy and, consequently, the speed of the initial transition. These include variations in the elasticity of substitution between clean and dirty energy, captured by the initial value σ_0, the damage coefficient γ, and the two preference parameters ρ and η. In general, these sensitivities are well known in the literature. When substitutability between clean and dirty energy is limited, optimal emissions never fall much below half of their initial level, leading to a much longer emissions half-life and increasing 2100 temperature. By contrast, greater substitutability facilitates deep decarbonisation and low optimal long-run temperatures. Similarly, a lower damage coefficient weakens climate ambition, resulting in both a slower transition and higher long-run temperature, while higher damages imply a much faster transition and lower temperature. With a high damage coefficient set at 4x its base value, which is in the direction of the recent results of Bilal and Känzig (2026) (our

SCC = \$556/tCO$_2$eq in 2025 in this case), the emissions half-life falls to just three years and temperature in 2100 is only 1.7 °C. This illustrates how recent, very high estimates of climate damages imply an extremely fast transition. The preference parameters operate through discounting (in this deterministic, globally aggregated setting, η affects outcomes solely through its influence on the effective discount rate). Lower discounting strengthens climate ambition and accelerates the initial transition, while higher discounting has the opposite effect.

The third type of relationship comprises parameter variations that do not have a strong effect on either the initial transition or long-run temperature. These include the technical change parameters, population growth g_L, and the rate of autonomous decarbonisation g_ψ. Varying these parameters produces only modest changes in both the emissions half-life and long-run temperature.

How robust is our overall finding of a fast transition? We find it is highly sensitive to assumptions about the climate response and damages, social preferences/discounting, and the substitutability of clean energy. But these sensitivities are well established in the literature as determinants of overall climate ambition and consequently the urgency of decarbonisation. Our result is moderately sensitive to assumptions about capital inertia and capital-energy substitution, with the emissions half-life varying by a few years, but overall we would argue this still constitutes a fast transition. Our result is robust to assumptions about clean technical change, growth and decarbonisation.

Decomposing transition inertia: expansion versus retirement frictions

The sensitivity analysis varies the general adjustment cost parameter χ and the stranding cost parameter χ_s jointly. To better understand the sources of transition inertia, we now vary these parameters independently and focus on the first two decades of the transition. This allows us to distinguish between frictions affecting the expansion of clean capital and frictions affecting the retirement of dirty capital. Figure 3 decomposes the resulting dynamics into capital accumulation (panel a), adjustment costs as a share of investment (b), aggregate energy supply (c), and emissions (d).

Panel a shows that doubling χ substantially slows the build-out of clean capital. Over the first decade, clean capital accumulation is markedly weaker than in the baseline, while dirty capital declines more slowly as a consequence of reduced clean substitution. By contrast, doubling the stranding cost parameter χ_s has little effect on clean capital accumulation and only modestly slows the initial decline of dirty capital.

Panel b clarifies the mechanism. When χ is doubled, the clean adjustment cost share ι_c/I_c rises significantly to above 35% and remains elevated, indicating that expansion of clean capacity is directly constrained by higher investment frictions. In contrast, the stranding

Table 1: Sensitivity analysis of key parameters.

Parameter/sensitivity		Range	P half-life (years)			T in 2100 (°C)		
			Low	Default	High	Low	Default	High
Capital inertia								
Adjustment costs	χ & χ_s	$[0, 2] \times$ base	8	10	14	2.0	2.1	2.2
Capital depreciation rate	δ	[2%, 6%]	16	10	7	2.1	2.1	2.1
Energy substitution								
Elasticity (capital–energy)	σ_f	[0.05, 1.01]	13	10	9	2.2	2.1	2.1
Elasticity (clean–dirty)	σ_0	[1.01, 3]	86	10	9	3.1	2.1	1.9
Climate and damages								
TCRE	ζ	$[4.2, 7.8]\times10^{-4}$	24	10	8	2.1	2.1	2.2
Damage coefficient	γ	$[0.01, 0.08]$	38	10	3	2.7	2.1	1.7
Preferences/discounting								
Utility discount rate	ρ	[0%, 2.5%]	4	10	45	1.7	2.1	2.9
Elasticity of marginal utility	η	[0.5, 2.2]	4	10	56	1.7	2.1	2.9
Clean technical change								
Speed of technical change	g_σ, ϱ	$[0, 2] \times$ base	8	10	10	2.6	2.1	2.1
Weight on exog/endog TC	ω	[0, 1]	10	10	11	2.4	2.1	2.1
Growth & decarbonisation								
Population growth	g_L	[0.23%, 0.37%]	10	10	10	2.2	2.1	2.2
Exogenous decarbonisation	g_ψ	[0.5%, 1.5%]	11	10	9	2.1	2.1	2.3

Notes: "Low" and "High" refer to the lower and upper bounds of the range shown.

cost share $|\Theta/I_d|$ changes little. When χ_s is doubled, the opposite pattern emerges: the stranding cost share rises in the early years, while the clean adjustment cost share remains close to its baseline path.

Panels c and d trace the aggregate consequences of these distinct frictions. Doubling χ produces a substantially deeper contraction in energy supply and a slower recovery thereafter. The resulting energy shortfall persists well into the second decade of the transition. Although tighter energy supply tends to dampen energy use and thereby emissions, this effect is more than offset by slower clean capital expansion and a more gradual decline of dirty capital. As a result, emissions remain persistently above the baseline path. By contrast, doubling χ_s has only modest effects on aggregate energy and emissions. Although dirty capital is retired somewhat more gradually and stranding costs are higher relative to disinvestment, clean expansion proceeds largely unchanged. As a result, total energy supply remains close to its baseline path and emissions deviate only slightly.

Taken together, these results imply that the critical bottleneck in the transition is how quickly clean capital can be deployed, not whether dirty capital is costly to retire. In the model, build-out frictions meaningfully constrain energy supply and delay decarbonisation, whereas retirement frictions have only modest macroeconomic effects when clean expansion is unconstrained.

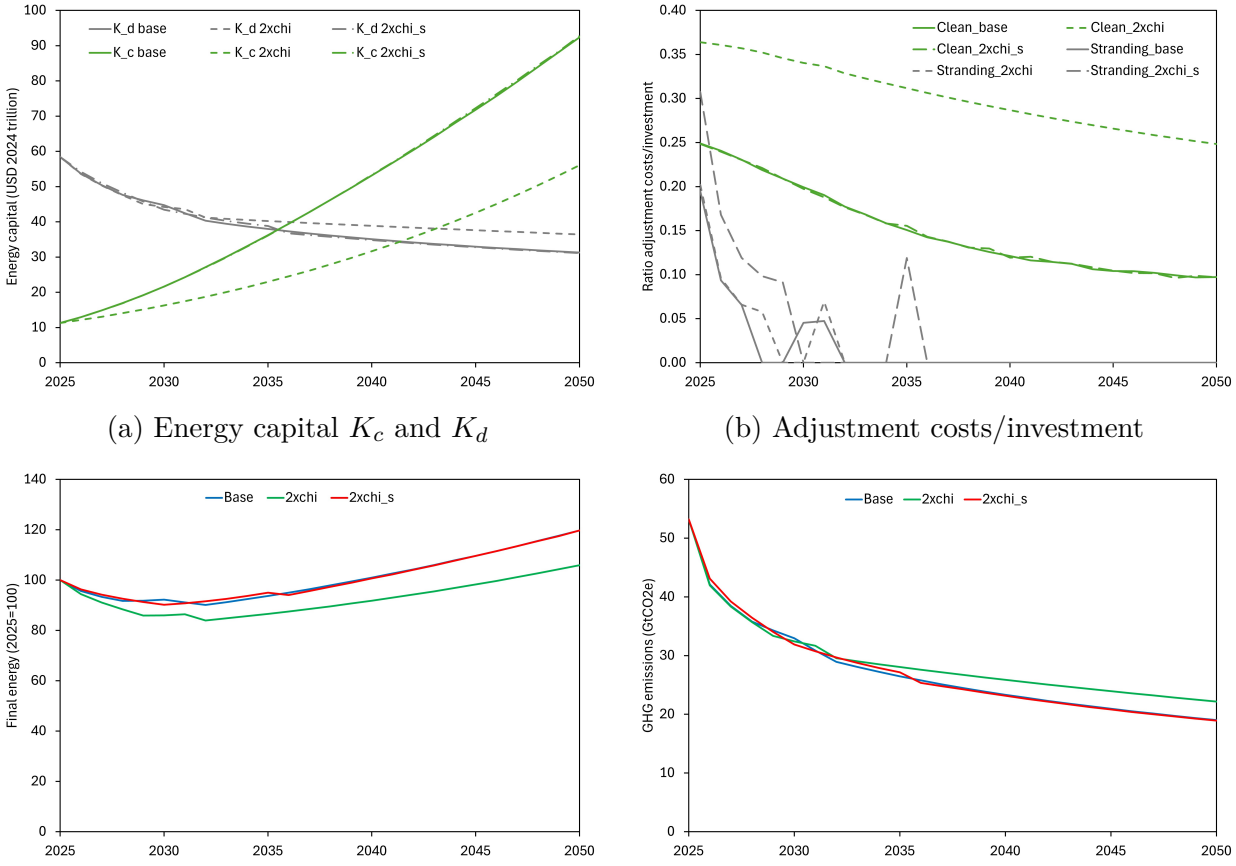

(a) Energy capital K_c and K_d

(b) Adjustment costs/investment

(c) Energy E

(d) Emissions P

Figure 3: Decomposing transition inertia: expansion versus retirement frictions. Panel (b) reports clean energy capital adjustment costs as a share of clean energy investment, and stranding adjustment costs as a share of dirty capital disinvestment.

The role of the dirty tail

A central feature of the model is heterogeneous dirty capital, and global data discussed above identify a tranche of highly emissions-intensive dirty capital: the dirty tail. Because these units generate disproportionately large emissions per unit of capital, their premature retirement delivers substantial emissions reductions per dollar of capital scrapped. To isolate the importance of this, we eliminate within-dirty heterogeneity by setting $\psi_2 = 0$. We then adjust ψ_3 so that, for given $K_d(0)$, initial emissions are unchanged and continue to match the observed level.

Figure 4 illustrates the consequences. Panel (a) shows the evolution of dirty capital, clean capital, and stranded capital, while panel (b) shows the path of emissions. Removing the dirty tail sharply reduces optimal premature retirement: the stock of stranded capital K_s peaks at only \$0.8tn, about five times less than the \$4 tn of stranded capital in the base

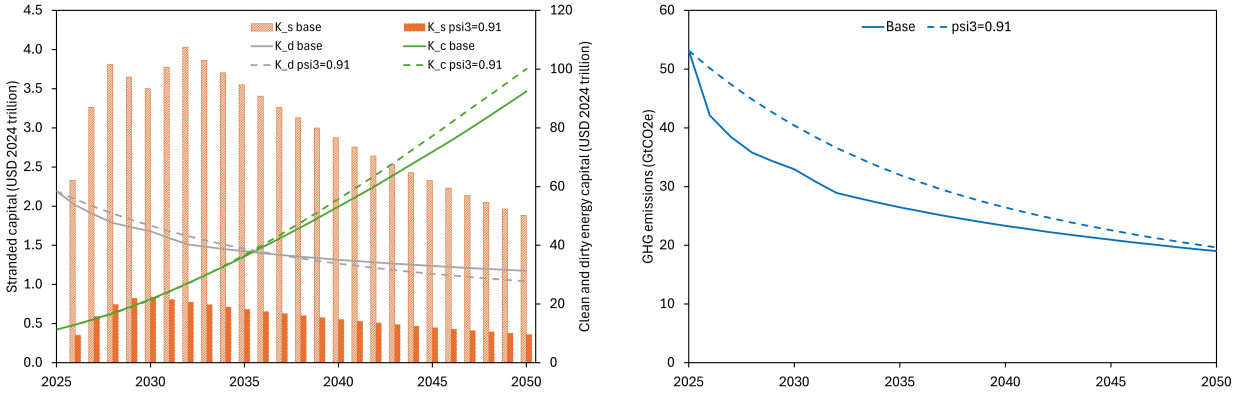

(a) Energy capital K_c, K_d and K_s

(b) Emissions P

Figure 4: Removing the dirty tail in the stranding function.

case. As a result, abatement shifts toward clean capital accumulation. Clean capital rises more strongly over time than in the baseline, while dirty capital declines more gradually in the early years. The loss of the high-return stranding margin leads to a more gradual decline in emissions over the first decade. By 2050, however, the emissions paths converge as clean capital eventually substitutes for dirty capital.

These results pin down the importance of heterogeneity in dirty capital for the timing of decarbonisation. The dirty tail accelerates early emissions reductions by making premature retirement an efficient abatement option. When this heterogeneity is removed, the transition relies more heavily on clean capital build-out and becomes more gradual in the near term, even though long-run outcomes are similar.

6 Stochastic results

In the preceding section, we showed that the optimal policy implies a rapid transition. We now examine how this result is affected by uncertainty. We focus on three questions: how uncertainty shapes the distribution of optimal transition paths; which sources of uncertainty matter for key outcomes; and how much the planner can adjust the transition in response to realised shocks, given capital inertia and adjustment costs.

The stochastic model is computationally challenging to solve, for several reasons. First, the model has six continuous state variables (the four capital stocks, cumulative emissions, and time) and three decision variables (investment in the three productive capital stocks). Second, the value function is highly non-linear and by design it has a kink at zero dirty investment. Third, we have geometric Brownian motions on the three productive capital stocks and on temperature. Recent work at the computational frontier of stochastic IAMs

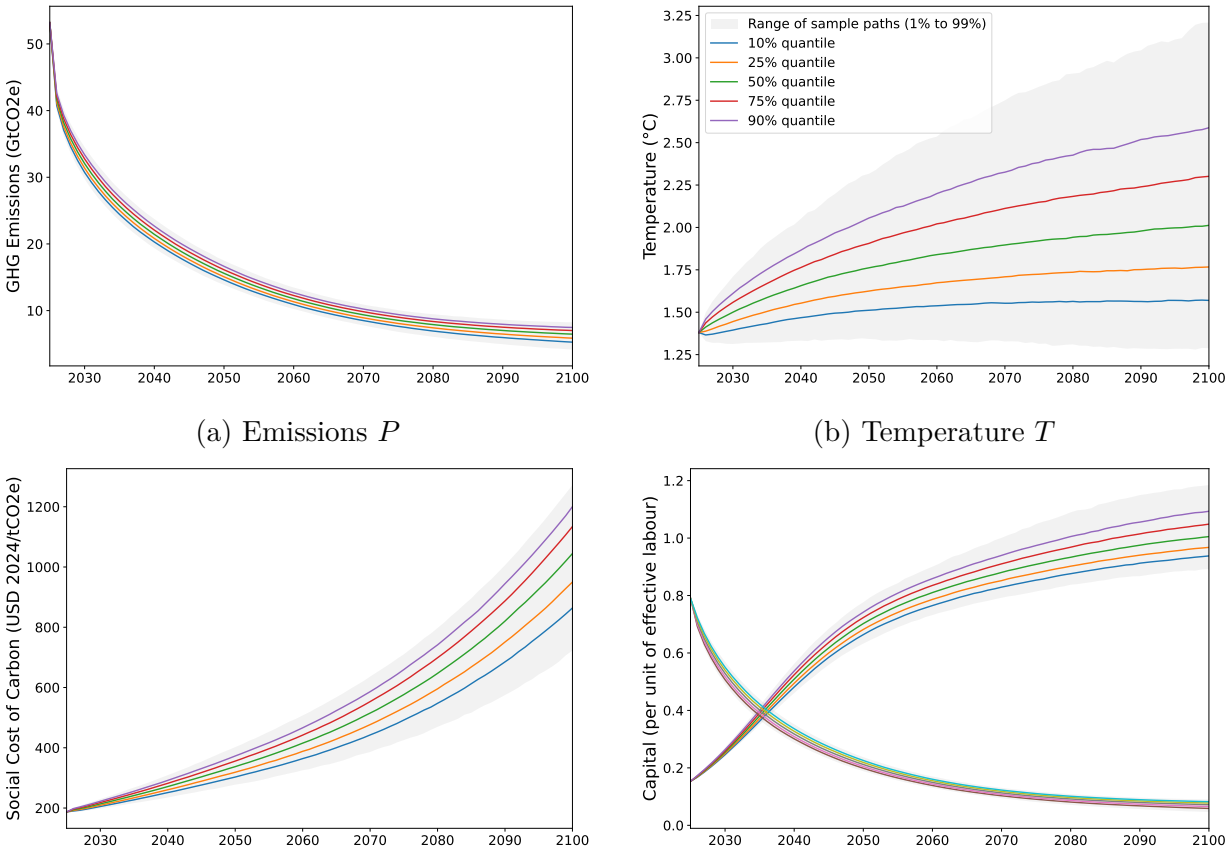

(a) Emissions P

(b) Temperature T

(c) SCC/Pigouvian carbon tax

(d) Energy capital k_c and k_d

Figure 5: Stochastic optimal transition paths.

includes Cai and Lontzek (2019); Friedl et al. (2023); Folini et al. (2024). We follow Azinovic et al. (2022) in using a deep neural net function approximator for the policy function based on the DEQN method. This enables us to handle the curse of dimensionality and the kink.

The neural net (policy function) maps states into optimal controls and shadow values, and is trained using simulated state trajectories rather than a fixed grid, focusing approximation accuracy on the relevant regions of the state space. Parameters are estimated by minimising the squared residuals of the model's optimality conditions using stochastic gradient descent. Appendix E provides further details.

The optimal transition under uncertainty

Figure 5 reports probability distributions for optimal emissions, temperature, the SCC/carbon tax, and clean and dirty capital, based on 10,000 Monte Carlo runs. Panels (a) and (d) show that emissions and capital paths exhibit relatively limited dispersion, particularly in the first decades. This reflects the fact that a large part of the transition is effectively predetermined

by the structure of the capital stock. Early on, emissions reductions are pinned down by the optimal retirement of highly emissions-intensive assets – the dirty tail – which is optimal across a wide range of states. Later in the transition, adjustment shifts toward the expansion of clean capital and the gradual replacement of lower-emissions dirty capital, but responses remain muted because it is costly to substitute away the last bit of dirty capital. As a result, both emissions and capital remain relatively insensitive to realised shocks throughout the transition.

By contrast, panels (b) and (c) show that temperature and the SCC are highly dispersed. Temperature in 2100 ranges from around 1.5 °C at the 10th percentile to 2.5 °C at the 90th percentile, while the SCC reaches \$280–360/tCO_2e in 2050. This reflects the fact that temperature remains highly uncertain even for similar emissions paths, due to large uncertainty in the climate response to emissions. Since damages are convex in temperature, this translates into large variation in marginal damages and hence in the SCC, even when emissions paths are similar.

These results imply an asymmetric response to realised shocks: when warming turns out to be worse than expected, the planner has limited scope to accelerate the transition, as key adjustment margins are already either exhausted or constrained. Conversely, when warming is less severe, emissions reductions remain substantial because the transition is largely shaped by adjustment constraints rather than marginal damages alone.

Decomposition of shocks

Figure 6 compares three specifications: (i) uncertainty in both temperature and capital, (ii) temperature uncertainty only, and (iii) capital uncertainty only. Temperature uncertainty has a clear effect on the SCC, particularly in the second half of the century (panel a), and leads to slightly lower temperatures on average (panel b), reflecting a modest insurance effect due to convex climate damages. The effect on emissions remains limited, consistent with the muted response of quantities documented above.

By contrast, capital uncertainty has little effect on the SCC or on temperature. In our calibration, this reflects the fact that the prudence effect – whereby consumption risk raises the marginal utility of future consumption – is largely offset by the risk premium on capital. In particular, because climate damages are proportional to output, the effective 'climate beta' is close to one, while intertemporal inequality aversion ($\eta = 1.35$) is of similar magnitude. In a model where $\eta = \beta$, a negative shock to GDP will increase marginal utility and decrease damages in equal proportion (Dietz et al., 2018; van den Bremer and van der Ploeg, 2021). As a result, the carbon price will not be affected by growth/macroeconomic uncertainty, as is the case in Golosov et al. (2014). Capital uncertainty nonetheless contributes to dispersion

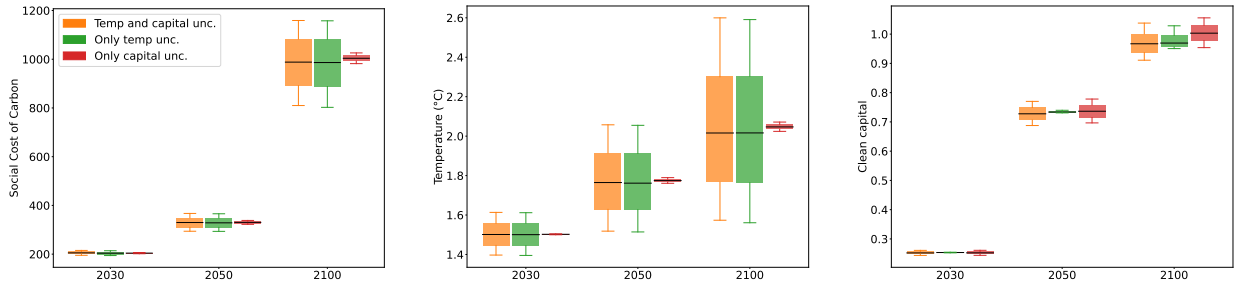

(a) SCC/Pigouvian carbon tax (b) Temperature T (c) Clean capital k_c

Figure 6: The effect of temperature and capital shocks on the mean SCC/Pigouvian tax, temperature and clean capital. The orange model has uncertainty on temperature and the three productive capital stocks, the green model only has temperature uncertainty and the red model only has capital uncertainty. The boxplots show percentiles P10, P25, P50, P75 and P90.

in clean capital, particularly in the second half of the century (panel c), where investment dynamics are more responsive to economic conditions.

Figure 7 reports impulse response functions (IRFs) for shocks to temperature, final-good capital, and clean capital. These IRFs can be interpreted as local measures of how much the planner can adjust the transition path in response to realised shocks.

The left column shows the response to a positive temperature shock of $+0.1\,^{\circ}\mathrm{C}$. Despite a noticeable increase in the carbon price (see Appendix F), emissions decline only modestly, by around $0.2\,\mathrm{GtCO_2e}$ in the very long run. As a result, the temperature response remains highly persistent. Note that in a model without emissions inertia, we would see an immediate reduction in emissions in year one followed by a gradual return to zero, both for emissions and for temperature. By contrast, in our model, adjustment costs lead to a 'sticky' emissions response, taking decades to adjust. And as temperature is the result of cumulative emissions, we obtain a 'super sticky' temperature.

The middle column shows the response to a positive shock to final-good capital. Higher capital raises energy demand, leading to an increase in both clean investment and emissions. The expansion of clean capital is limited due to adjustment costs, so part of the additional energy demand is met by decelerating stranding of dirty capital. This highlights the role of energy as a bottleneck: macroeconomic expansions translate into higher emissions when clean capacity cannot adjust quickly enough.

The right column shows the response to a positive shock to clean capital. The immediate effect is a reallocation away from clean investment, but this is limited because clean investment is already close to its maximum expansion path. Over time, higher clean capacity supports higher output and energy demand, leading to a partial rebound in emissions. The

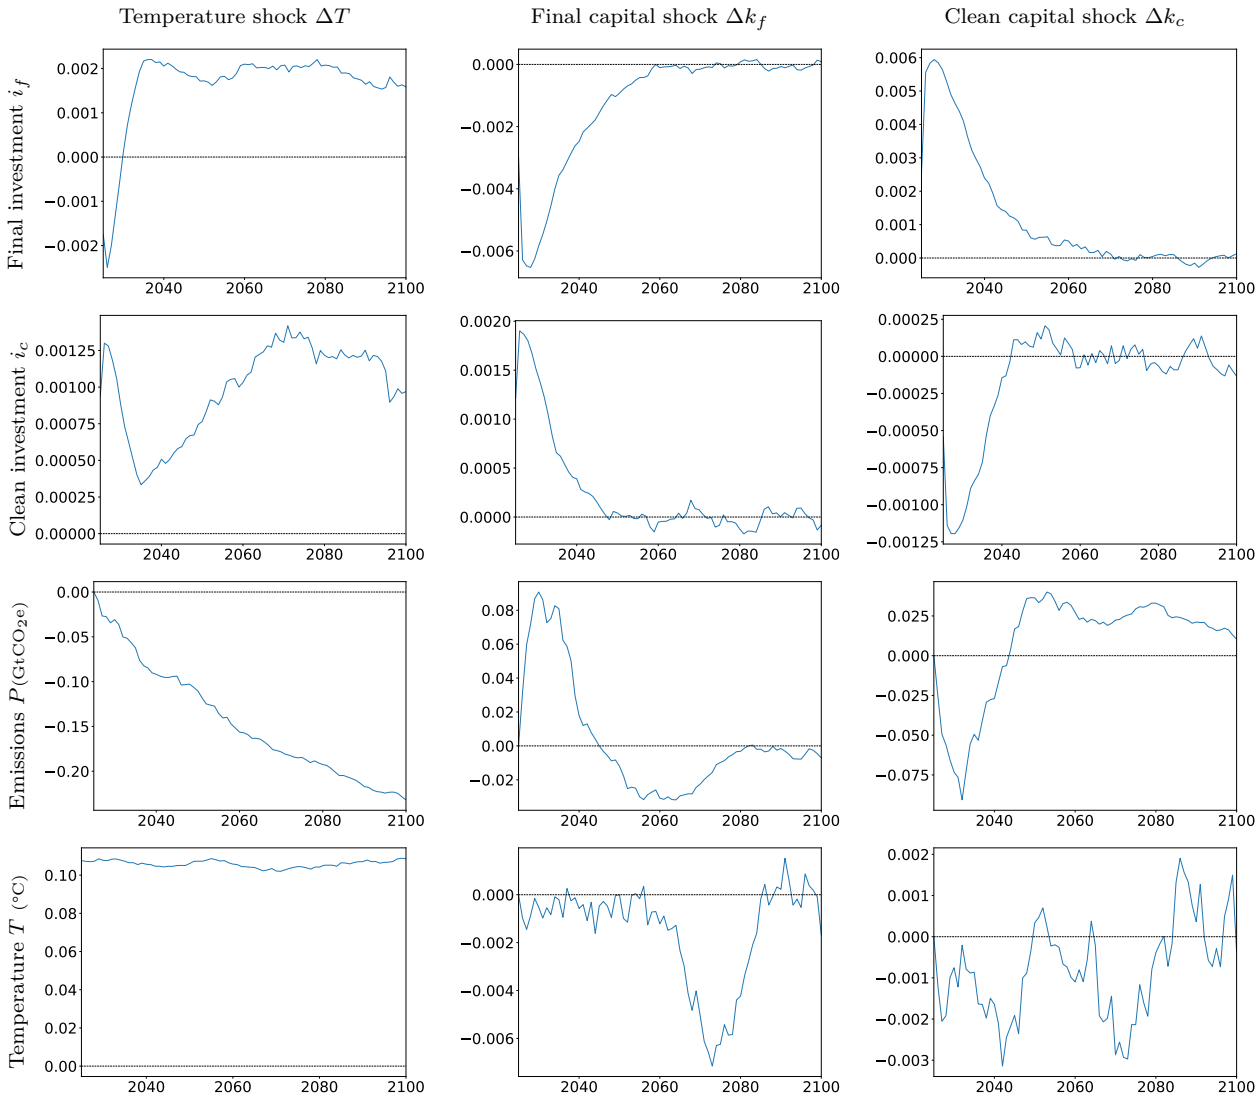

Figure 7: Impulse response functions. Columns correspond to a positive shock on temperature, final capital and clean capital respectively. Rows show the response of final investment, clean investment, emissions and temperature. Investment is expressed per unit of effective labour. The size of the shocks corresponds to three standard deviations. The response is the mean difference compared to the baseline for 10,000 Monte Carlo runs.

shock is absorbed only gradually, reflecting both capital persistence and the central role of energy constraints in shaping investment dynamics.

Across all shocks, a common pattern emerges: prices adjust more than quantities. While the planner responds actively through the carbon price, emissions respond only modestly, especially in the short run. This reflects the fact that the transition is constrained along two margins. Early on, emissions reductions are effectively predetermined by the retirement of the dirty tail. Thereafter, adjustment is governed by the costs of expanding clean capital

and substituting away from the last remaining dirty capital. As a result, the scope for *ex post* adjustment of the transition path is limited.

7 Discussion

This paper studies optimal climate policy when the low-carbon transition is constrained by frictions in retiring dirty assets and scaling clean capital. We embed empirically disciplined heterogeneous dirty capital and adjustment costs in a general equilibrium model. We show that optimal policy entails rapid near-term decarbonisation. The key mechanism is a 'dirty tail' of highly emissions-intensive assets that can be retired at relatively low cost, delivering a sharp initial decline in emissions. As the transition unfolds, the margin of adjustment shifts from early retirement of the dirty tail to the expansion of clean capital. Accordingly, stranding costs are not strongly binding, whereas limits to clean investment constitute the central bottleneck. Extending the analysis to a stochastic setting reinforces this picture: uncertainty primarily affects temperatures and marginal damages, while optimal emissions and investment paths remain comparatively insensitive, reflecting the limited flexibility created by capital inertia and the front-loaded optimal retirement of the dirty tail.

Here we benchmark our quantitative results against the existing literature as a consistency check, showing that our model delivers (i) carbon prices, (ii) abatement costs, (iii) energy demand, (iv) stranded dirty capital and (v) risk premia that are broadly in line with recent work, and that any differences can be traced to identifiable features of our model:

(i) We estimate a Pigouvian carbon tax/SCC of \$181/tCO_2e in 2025, rising to \$331/tCO_2e by 2050. These values are consistent with recent work reporting substantially higher SCC values, as discussed in the introduction. The consistency comes primarily from rapid climate dynamics (Dietz et al., 2021) and high damages (Howard and Sterner, 2017) that recent studies share.

(ii) We compare the MAC curve implied by our model to the range of MACs in the IPCC AR6 and NGFS scenario databases (IPCC, 2022; NGFS, 2024). Our implied MAC curve lies within the interquartile range of other IAMs (Figure A3 in Appendix F.2), except near current emissions levels. The very low initial MACs in the IPCC/NGFS models are partly explained by those models solving a cost-effectiveness/Hotelling-type problem under relatively high discount rates of 4-5% and ignoring macro-economic adjustment costs. The consistency of our estimates with other IAMs is reassuring given that we do not explicitly target them in calibration.

(iii) Our model implies a temporary decline in the aggregate energy input E of around 10% over the first decade of the transition. While E is not directly comparable to standard

IAM energy aggregates, this pattern can be benchmarked against global final energy demand in the IPCC AR6 scenarios (IPCC, 2022). Using reported ranges for 2020 and 2030, one can construct conservative outer bounds for the implied change in energy demand of roughly 80-130 on a 2020=100 index. Our projected decline lies comfortably within the range.

(iv) Our estimates of stranded assets are also broadly in line with the literature, once differences in definitions are taken into account. The stock of stranded dirty capital peaks at around \$4 tn (about 3.6% of initial GDP). Direct comparison is difficult, since stranding is defined differently, often measured in physical rather than monetary units, and typically focuses on narrower sectors. von Dulong (2023) and Fofrich Navarro et al. (2026) estimate stranded fossil power assets of \$500 bn in a 2 °C scenario, while Edwards et al. (2022) estimate stranded coal assets of \$0.6-1 tn alone. Using a different definition of capital stranding, Burda et al. (2025) estimate losses equal to 1.46% of GDP under ambitious climate policy. Since our analysis covers the entire productive system and identifies the optimal stranding path endogenously, our estimates are naturally larger. Nevertheless, they remain broadly consistent with the existing literature.

(v) We find a small effect of mean-preserving temperature uncertainty on the SCC. Other studies investigating climate uncertainties and tipping points find similar results (Cai and Lontzek, 2019; Crost and Traeger, 2014; van den Bremer and van der Ploeg, 2021).[13] First, although cumulative temperature shocks lead to large uncertainty over a century, the yearly effects are modest compared to other shocks. A one standard deviation shock on temperature leads to a 0.1% reduction in GDP, which is less than a tenth of the observed standard deviation of world GDP. Second, on an optimal trajectory, shocks are less harmful, as marginal damages are balanced with abatement costs. On an non-optimal trajectory, however, the risk premium on the carbon price can be substantial.

The concentration of emissions in a relatively small set of highly emissions-intensive assets – the dirty tail – implies that effective mitigation can, in principle, be achieved by targeting a limited subset of the capital stock. In our model, a Pigouvian carbon price implements this outcome directly. Experience shows, however, that a global carbon price is difficult to implement at the required scale and credibility.

This motivates an inquiry into second-best instruments that could more directly target the dirty tail. To start with, one could relax the assumption of a globally harmonised

[13] Cai and Lontzek (2019) find that increasing the RRA from 2 to 5 increases the SCC by a mere 3% (Fig. 7 therein). In Hambel et al. (2021), the effect is 3% for the Dell et al. (2012) damages calibration and $\eta = 1$ (Table 6 therein). The effect is barely visible in Crost and Traeger (2014), SM Fig 5.

carbon price and consider differentiated prices, presumably lower in developing countries. While inefficient from a global perspective, this may still be sufficient to incentivise early retirement of the dirty tail given the low costs involved. Moving beyond Pigouvian carbon prices, a range of policy options exists. In a national context, these include government-mandated phase-out schedules, reverse auctions, performance or rate-based standards, and feebate schemes. In an international context, they include targeted finance and technology transfer to support early retirement of high-emissions assets in developing countries. Existing initiatives, such as Just Energy Transition Partnerships under the Paris Agreement and coal transition vehicles, provide examples of this approach.

Our analysis also highlights the importance of easing bottlenecks on the clean side of the transition. Constraints on the rate of clean capital expansion – arising from manufacturing capacity, supply chains, permitting, and infrastructure – materially affect the pace of emissions reductions, especially after the dirty tail is eliminated. This points to the importance of policies that relax these constraints, including support for manufacturing capacity, permitting reform, grid infrastructure, and supply chains for critical materials.

An effective second-best policy mix therefore combines targeted measures to accelerate the retirement of the dirtiest assets with broad-based policies to expand clean capacity. Early stranding of the dirty tail delivers large emissions reductions at relatively low cost, while accelerating clean investment addresses the binding constraint on the speed of the transition. Designing policies that achieve this combination in practice remains an important challenge.

Acknowledgements

We acknowledge funding from the National Bank of Belgium, the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (ref. 853050 - SMOOTH), the European Union's Horizon Europe Research and Innovation Programme (ref. 101170509 - BELIEFS), the ESRC Centre for Climate Change Economics and Policy (ref. ES/R009708/1), and the Grantham Foundation for the Protection of the Environment. The authors have no conflicts of interest or financial interests to disclose.

References

Acemoglu, D., Aghion, P., Bursztyn, L., Hemous, D., 2012. The Environment and Directed Technical Change. American Economic Review 102, 131–166. doi:10.1257/aer.102.1.131.

Annicchiarico, B., Carattini, S., Fischer, C., Heutel, G., 2022. Business Cycles and Environ-

mental Policy: A Primer. Environmental and Energy Policy and the Economy 3, 221–253. doi:10.1086/717222.

Asker, J., Collard-Wexler, A., De Loecker, J., 2014. Dynamic Inputs and Resource (Mis)Allocation. Journal of Political Economy 122, 1013–1063. doi:10.1086/677072.

Azinovic, M., Gaegauf, L., Scheidegger, S., 2022. Deep Equilibrium Nets. International Economic Review 63, 1471–1525. doi:10.1111/iere.12575.

Baldwin, E., Cai, Y., Kuralbayeva, K., 2020. To build or not to build? Capital stocks and climate policy. Journal of Environmental Economics and Management 100, 102235. doi:10.1016/j.jeem.2019.05.001.

Barnett, M., Brock, W., Hansen, L.P., 2020. Pricing Uncertainty Induced by Climate Change. The Review of Financial Studies 33, 1024–1066. doi:10.1093/rfs/hhz144.

Barrage, L., 2018. Be careful what you calibrate for: Social discounting in general equilibrium. Journal of Public Economics 160, 33–49. doi:10.1016/j.jpubeco.2018.02.012.

Barrage, L., 2020. Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy. The Review of Economic Studies 87, 1–39. doi:10.1093/restud/rdz055.

Barrage, L., Nordhaus, W., 2024. Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. Proceedings of the National Academy of Sciences 121, e2312030121.

Barro, R.J., 2009. Rare disasters, asset prices, and welfare costs. American Economic Review 99, 243–64.

BEIS, 2024. Digest of UK Energy Statistics (DUKES) 2024. Technical Report. Department for Energy Security and Net Zero.

Bennett, L.L., Skjoeldebrand, R., 1986. Worldwide nuclear power status and trends. IAEA Bulletin 28.

Bilal, A., Känzig, D.R., 2026. The Macroeconomic Impact of Climate Change: Global Versus Local Temperature. The Quarterly Journal of Economics , qjag011doi:10.1093/qje/qjag011.

Burda, M.C., Goeth, A.M., Zessner-Spitzenberg, L., 2025. Capital Adjustment Costs and Stranded Assets in an Optimal Energy Transition. Discussion Paper 18356. IZA Institute of Labor Economics.

Cai, Y., 2019. Computational methods in environmental and resource economics. Annual Review of Resource Economics 11, 59–82. doi:10.1146/annurev-resource-100518-093841.

Cai, Y., Lontzek, T.S., 2019. The social cost of carbon with economic and climate risks. Journal of Political Economy 127, 2684–2734. doi:10.1086/701890.

Coppens, L., Dietz, S., Venmans, F., 2025. Optimal climate policy under exogenous and endogenous technical change: Making sense of the different approaches. Journal of Environmental Economics and Management 133, 103216. doi:10.1016/j.jeem.2025.103216.

Crippa, M., Guizzardi, D., Pagani, F., Banja, M., Muntean, M., et al., 2025. GHG Emissions of All World Countries – 2025 Report. Report JRC143227. Publications Office of the European Union. Luxembourg. doi:10.2760/9816914.

Crost, B., Traeger, C.P., 2014. Optimal CO2 mitigation under damage risk valuation. Nature Climate Change 4, 631–636.

David, J.M., Venkateswaran, V., 2019. The Sources of Capital Misallocation. American Economic Review 109, 2531–2567. doi:10.1257/aer.20180336.

Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature shocks and economic growth: Evidence from the last half century. American Economic Journal: Macroeconomics 4, 66–95.

Dietz, S., Gollier, C., Kessler, L., 2018. The climate beta. Journal of Environmental Economics and Management 87, 258–274.

Dietz, S., Rising, J., Stoerk, T., Wagner, G., 2021. Economic impacts of tipping points in the climate system. Proceedings of the National Academy of Sciences 118.

Dietz, S., Venmans, F., 2019. Cumulative carbon emissions and economic policy: In search of general principles. Journal of Environmental Economics and Management 96, 108–129. doi:10.1016/j.jeem.2019.04.003.

Drupp, M.A., Freeman, M.C., Groom, B., Nesje, F., 2018. Discounting Disentangled. American Economic Journal: Economic Policy 10, 109–134. doi:10.1257/pol.20160240.

Duffie, D., Epstein, L.G., 1992. Asset Pricing with Stochastic Differential Utility. The Review of Financial Studies 5, 411–436. doi:10.1093/rfs/5.3.411.

Edwards, M.R., Cui, R., Bindl, M., et al., 2022. Quantifying the regional stranded asset risks from new coal plants under 1.5°C. Environmental Research Letters 17, 024029. doi:10.1088/1748-9326/ac4ec2.

EIA, 2023. Planned Retirements of U.S. Coal-Fired Electric-Generating Capacity to Increase in 2025. In-Brief Analysis. U.S. Energy Information Administration.

EIA, 2025a. Cost and Performance Characteristics of New Generating Technologies, Annual Energy Outlook 2025. Technical Report. U.S. Energy Information Administration. Washington, DC.

EIA, 2025b. Electric Generator Report (EIA-860) — Capacity Additions and Retirements. Technical Report. U.S. Energy Information Administration.

EPA, 2023. Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances. Technical Report. Environmental Protection Agency. Washington, DC.

Eurostat, 2015. Manual for Air Emissions Accounts. Technical Report. European Union. Luxembourg.

Fofrich Navarro, R., Liebermann, L., Moore, F.C., Shearer, C., Davis, S.J., 2026. Ownership of power plants stranded by climate mitigation. Nature Sustainability 9, 328–336. doi:10.1038/s41893-025-01707-5.

Folini, D., Friedl, A., Kübler, F., Scheidegger, S., 2024. The Climate in Climate Economics. The Review of Economic Studies , rdae011doi:10.1093/restud/rdae011.

Fried, S., 2018. Climate policy and innovation: A quantitative macroeconomic analysis. American Economic Journal: Macroeconomics 10, 90–118. doi:10.1257/mac.20150289.

Friedl, A., Kübler, F., Scheidegger, S., Usui, T., 2023. Deep Uncertainty Quantification: With an Application to Integrated Assessment Models. Technical Report.

GEM, 2026. Global Energy Monitor.

Gerlagh, R., Liski, M., 2018. Carbon prices for the next hundred years. The Economic Journal 128, 728–757.

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica : journal of the Econometric Society 82, 41–88. doi:10.3982/ECTA10217.

Goulder, L.H., Mathai, K., 2000. Optimal CO2 abatement in the presence of induced technological change. Journal of Environmental Economics and Management 39, 1–38.

Hambel, C., Kraft, H., Schwartz, E., 2021. Optimal carbon abatement in a stochastic equilibrium model with climate change. European Economic Review 132, 103642.

Hambel, C., Kraft, H., van der Ploeg, F., 2024. Asset Diversification Versus Climate Action. International Economic Review 65, 1323–1355. doi:10.1111/iere.12691.

Hassler, J., Krusell, P., Olovsson, C., 2021. Directed Technical Change as a Response to Natural Resource Scarcity. Journal of Political Economy 129, 3039–3072. doi:10.1086/715849.

Hochmuth, P., Krusell, P., Mitman, K., 2026. Distributional Consequences of Becoming Climate-Neutral*. The Economic Journal , ueag004doi:10.1093/ej/ueag004.

Howard, P.H., Sterner, T., 2017. Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates. Environmental and Resource Economics 68, 197–225. doi:10.1007/s10640-017-0166-z.

IEA, 2020. Energy Technology Perspectives 2020. Technical Report. International Energy Agency. Paris.

IEA, 2021. Net Zero by 2050 - A Roadmap for the Global Energy Sector. Technical Report. International Energy Agency.

IEA, 2025. World Energy Outlook 2025. Technical Report. International Energy Agency. Paris.

IPCC, 2006. 2006 IPCC Guidelines for National Greenhouse Gas Inventories. Technical Report. Intergovernmental Panel on Climate Change.

IPCC, 2014. Annex III: Technology-specific cost and performance parameters, in: Climate Change 2014: Mitigation of Climate Change. Cambridge University Press. volume Working Group III - Intergovernmental Panel on Climate Change.

IPCC, 2021. Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA.

IPCC, 2022. AR6 scenarios database. doi:10.5281/zenodo.5886911.

Iverson, T., Karp, L., 2021. Carbon Taxes and Climate Commitment with Non-constant Time Preference. The Review of Economic Studies 88, 764–799. doi:10.1093/restud/rdaa048.

Jackson, S.B., (Kelvin) Liu, X., Cecchini, M., 2009. Economic consequences of firms' depreciation method choice: Evidence from capital investments. Journal of Accounting and Economics 48, 54–68. doi:10.1016/j.jacceco.2009.06.001.

Jindal, A., Shrimali, G., 2022. Cost–benefit analysis of coal plant repurposing in developing countries: A case study of India. Energy Policy 164, 112911. doi:10.1016/j.enpol.2022.112911.

Kermani, A., Ma, Y., 2023. Asset Specificity of Nonfinancial Firms. The Quarterly Journal of Economics 138, 205–264. doi:10.1093/qje/qjac030.

Lanteri, A., 2018. The Market for Used Capital: Endogenous Irreversibility and Reallocation over the Business Cycle. American Economic Review 108, 2383–2419. doi:10.1257/aer.20160131.

Lemoine, D., Traeger, C., 2014. Watch your step: Optimal policy in a tipping climate. American Economic Journal: Economic Policy 6, 137–66.

Moore, F.C., Drupp, M.A., Rising, J., Dietz, S., Rudik, I., Wagner, G., 2024. Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties. Proceedings of the National Academy of Sciences 121, e2410733121.

Mullen, K.M., Ardia, D., Gil, D.L., Windover, D., Cline, J., 2011. DEoptim: An R Package for Global Optimization by Differential Evolution. Journal of Statistical Software 40, 1–26. doi:10.18637/jss.v040.i06.

Nan, Y., Moseley, M., 2011. The Expansion of China's Generation Capacity. PPP Insights 1/1. Public–Private Infrastructure Advisory Facility.

NGFS, 2024. NGFS climate scenarios data set. doi:10.5281/zenodo.13989530.

Nordhaus, W.D., 1991. To slow or not to slow: The economics of the greenhouse effect. Economic Journal 101, 920–937.

Olovsson, C., Vestin, D., 2023. Greenflation? Working Paper Series .

Papageorgiou, C., Saam, M., Schulte, P., 2017. Substitution between Clean and Dirty Energy Inputs: A Macroeconomic Perspective. The Review of Economics and Statistics 99, 281–290. doi:10.1162/REST_a_00592.

Pindyck, R.S., Wang, N., 2013. The economic and policy consequences of catastrophes. American Economic Journal: Economic Policy 5, 306–39.

Raimi, D., 2017. Decommissioning US Power Plants: Decisions, Costs, and Key Issues.

Technical Report. Resources for the Future. Washington, DC.

Raimi, D., 2020. Decommissioning Power Plants: An Overview of the Financial, Regulatory, and Technical Challenges. Technical Report. Resources for the Future. Washington, DC.

Rozenberg, J., Vogt-Schilb, A., Hallegatte, S., 2020. Instrument choice and stranded assets in the transition to clean capital. Journal of Environmental Economics and Management 100, 102183. doi:10.1016/j.jeem.2018.10.005.

Sahuc, J.G., Smets, F., Vermandel, G., 2024. The New Keynesian Climate Model. SSRN Electronic Journal doi:10.2139/ssrn.5035207.

Taalbi, J., 2020. Evolution and structure of technological systems - An innovation output network. Research Policy 49, 104010. doi:10.1016/j.respol.2020.104010.

Tong, D., Zhang, Q., Zheng, Y., Caldeira, K., Shearer, C., Hong, C., Qin, Y., Davis, S.J., 2019. Committed emissions from existing energy infrastructure jeopardize 1.5°C climate target. Nature 572, 373–377. doi:10.1038/s41586-019-1364-3.

Traeger, C.P., 2023. ACE—Analytic Climate Economy. American Economic Journal: Economic Policy 15, 372–406. doi:10.1257/pol.20210297.

UN, 2024. World Population Prospects 2024: Summary of Results. Technical Report UN DESA/POP/2024/TR/NO. 9. United Nations Department of Economic and Social Affairs, Population Division.

UNFCCC, 2026. Greenhouse Gas Inventory Data - Detailed data by Party.

van den Bremer, T.S., van der Ploeg, F., 2021. The risk-adjusted carbon price. American Economic Review 111, 2782–2810.

van der Ploeg, F., Rezai, A., 2020a. The risk of policy tipping and stranded carbon assets. Journal of Environmental Economics and Management 100, 102258. doi:10.1016/j.jeem.2019.102258.

van der Ploeg, F., Rezai, A., 2020b. Stranded assets in the transition to a carbon-free economy. Annual Review of Resource Economics 12, 281–298.

von Dulong, A., 2023. Concentration of asset owners exposed to power sector stranded assets may trigger climate policy resistance. Nature Communications 14, 6442. doi:10.1038/s41467-023-42031-w.

WMO, 2025. State of the Global Climate 2024. Technical Report. World Meteorological Organization. Geneva.

World Bank, 2026. World Development Indicators.

WSB, 2019. Final Report. Technical Report. Commission on Growth, Structural Change and Employment - Federal Ministry for Economic Affairs and Climate Action. Berlin.

Appendix A Social planner solution of the deterministic model

This appendix finds the social planner's solution of the deterministic model.

The objective function for the deterministic model is

$$\int_0^\infty e^{-\rho t} L \frac{(C/L)^{1-\eta}}{1-\eta} dt. \tag{A1}$$

Expressed per unit of effective labour, the objective can be rewritten as

$$\int_0^\infty \exp(-\rho t)\, A^{1-\eta} L \frac{\left(\frac{C}{AL}\right)^{1-\eta}}{1-\eta}\, dt = \int_0^\infty \exp\left(-(\overbrace{\rho - g_L + (\eta-1)g}^{r-g})t\right)\, A(0)^{1-\eta} L(0) \frac{\left(\frac{C}{AL}\right)^{1-\eta}}{1-\eta}\, dt, \tag{A2}$$

where we define the effective discount rate as $r-g$ with $r = \rho + \eta g - g_L$ being the population-adjusted Ramsey rule. Production, energy and consumption follow equations 1, 2 and 18.

The Hamiltonian of the model is

$$\begin{aligned} H = & \frac{c^{1-\eta}}{1-\eta} \\ & + \lambda^{k_f} \left(i_f - \chi \frac{i_f^2}{k_f} - (\delta + g + g_L) k_f \right) \\ & + \lambda^{k_c} \left(i_c - \chi \frac{i_c^2}{k_c} - (\delta + g + g_L) k_c \right) \\ & + \lambda^{k_d} \left(i_d - \mathbf{1}_{i_d>0} \chi \frac{i_d^2}{k_d} - (\delta + g + g_L) k_d \right) \\ & + \lambda^{k_s} \left(-(\delta + g + g_L) k_s + \mathbf{1}_{i_d<0}(-i_d) \right) \\ & + \lambda^S P(k_d, k_s), \end{aligned}$$

where each shadow price is multiplied by its respective equation of motion (ex: $\dot{S} = P(k_d, k_s)$). Defining $a_j \equiv i_j/k_j$, the first-order condition (FOC) with respect to clean capital is

$$c^{-\eta} = \lambda^{k_c} \left(1 - 2\chi a_c\right), \tag{A3}$$

and the equivalent applies to final capital.

For dirty investments, we have three FOCs associated with three regimes: negative in-

vestment, no investment, and positive investment:[14]

$$i_d < 0: \quad 2\chi_s a_d c^{-\eta} + \lambda^{k_s} - \lambda^{k_d} = 0, \tag{A4}$$

$$i_d = 0: \quad \lambda^{k_s} \leq \lambda^{k_d}, \tag{A5}$$

$$i_d > 0: \quad c^{-\eta} = \lambda^{k_d} \left(1 - 2\chi a_d\right). \tag{A6}$$

In the FOC of the negative investment regime, the first term involving i_d is negative and captures the adjustment cost of stranding. The shadow price of stranding capital, λ^{k_s}, reflects the welfare gain from replacing a unit of dirty capital with the highest emissions intensity by one with the lowest emissions intensity. Over the stranding period, λ^{k_s} declines as the emissions intensity of the dirtiest capital decreases.

The shadow price of dirty capital, λ^{k_d}, reflects two opposing effects: the productivity of dirty capital as well as the marginal emissions associated with dirty capital, P_{k_d} (which reflect the emissions intensity of the least-polluting remaining units when dirty capital is heterogeneous). As a result, the initial sign of λ^{k_d} is ambiguous. During the stranding phase, the decline in k_d raises the marginal productivity of dirty capital, so λ^{k_d} increases up to the point where it equals λ^{k_s} when $i_d = 0$.

The FOC for zero investment is a regularity condition. If it were violated, it would be optimal to simultaneously invest in and disinvest from dirty capital. In that case, the problem would not be concave in i_d, and multiple local optima could arise. The FOC for positive investment is the same as for the other capital stocks.

The current-value costate equations are

$$\dot{\lambda}^{k_f} = (r - g)\lambda^{k_f} - c^{-\eta} y_{k_f} + \lambda^{k_f} \left(\delta + g + g_L - \chi a_f^2\right) \tag{A7}$$

$$\dot{\lambda}^{k_c} = (r - g)\lambda^{k_c} - c^{-\eta} y_{k_c} + \lambda^{k_c} \left(\delta + g + g_L - \chi a_c^2\right) \tag{A8}$$

$$\dot{\lambda}^{k_d} = (r - g)\lambda^{k_d} - c^{-\eta} \left(y_{k_d} + \mathbf{1}_{i_d<0} \chi_s a_d^2\right) + \lambda^{k_d} \left(\delta + g + g_L - \mathbf{1}_{i_d>0} \chi a_d^2\right) - \lambda^S P_{k_d} \tag{A9}$$

[14] In the numerical applications, we use a softplus function $\mathrm{sp}_\varepsilon(i_d) = \ln(1 + \exp(\varepsilon i d))/\varepsilon$ to smooth the kink. Its derivative is the sigmoid function $\varsigma_\varepsilon(i_d) = \mathrm{sp}'_\varepsilon(i_d) = \frac{1}{1+\exp(-\varepsilon i_d)}$. This leads to a single FOC:

$$\left(\varsigma_\varepsilon(i_d) + \mathbf{1}_{i_d<0}\, 2\chi_s a_d\right) c^{-\eta} + \left(1 - \varsigma_\varepsilon(i_d)\right) \lambda^{k_s} = \lambda^{k_d} \left(1 - \mathbf{1}_{i_d>0}\, 2\chi a_d\right).$$

Note that we smooth the payoff kink but retain regime-dependent adjustment costs.

$$\dot{\lambda}^{k_s} = (r + \delta + g_L)\,\lambda^{k_s} - \lambda^S P_{k_s} \tag{A10}$$

$$\dot{\lambda}^S = (r-g)\lambda^S - c^{-\eta}\left(-y\gamma\zeta^2 S - y_e e_\epsilon \epsilon_M \frac{1}{M(0)}\right) \tag{A11}$$

Integrating Equation (A11) shows that the shadow price of cumulative emissions corresponds to the sum of the social cost of carbon and the marginal social benefit of technological learning,

$$\lambda^S(t) = -\int_t^\infty e^{-(r-g)(\tau-t)} c(\tau)^{-\eta}\left[y(\tau)\gamma\zeta^2 S(\tau) + y_e(\tau)e_\epsilon(\tau)\epsilon_M \frac{1}{M(0)}\right] d\tau. \tag{A12}$$

At the optimum, the shadow price of emissions, λ^S, determines the marginal abatement cost (MAC). To obtain an expression for the MAC, we compare the optimality conditions for clean and dirty capital.

Start by rewriting the costate equation for k_c as follows:

$$r + g_L = \frac{\dot{\lambda}^{k_c}}{\lambda^{k_c}} + \frac{c^{-\eta}}{\lambda^{k_c}} y_{k_c} - \delta + \chi a_c^2. \tag{A13}$$

This equation is a Ramsey-type Euler equation with adjustment costs: the left-hand side is the discount rate adjusted for population growth, while the right-hand side gives the 4 components of the return on capital: capital gains, dividend yield (multiplied by Tobin's q), depreciation and adjustment costs (increasing k_c leads to lower adjustment costs).

Combining this with the FOC for clean investment, and defining $g_j^I \equiv \dot{i}_j/i_j$ and $g_j^K \equiv \dot{k}_j/k_j$, yields[15]

$$r + g_L + \eta\frac{\dot{c}}{c} + \delta = (1 - 2\chi a_c)\, y_{k_c} + \chi a_c^2 + \frac{g_c^I - g_c^K}{\frac{k_c}{2\chi i_c} - 1}. \tag{A14}$$

Applying the same steps to dirty capital (for non-negative investment) gives

$$r + g_L + \eta\frac{\dot{c}}{c} + \delta = (1 - 2\chi a_d)\, y_{k_d} + \chi a_d^2 + \frac{g_d^I - g_d^K}{\frac{k_d}{2\chi i_d} - 1} + \frac{\lambda^S}{\lambda^{k_d}} P_{k_d}. \tag{A15}$$

Subtracting the clean-capital equation from the dirty-capital equation eliminates the common left-hand side and yields an expression linking the shadow price of emissions to the wedge between dirty and clean returns.

[15]Note that $-\eta\dot{c}/c$ is close to zero and that r includes the growth of A.

Defining the MAC in per-effective-labour consumption units as

$$\text{MAC} \stackrel{\text{def}}{=} -\lambda^S c^\eta, \tag{A16}$$

we obtain

$$\text{MAC} = \frac{(1-2\chi a_d)\, y_{k_d} - (1-2\chi a_c)\, y_{k_c} + \frac{g_d^I - g_d^K}{\frac{k_d}{2\chi i_d}-1} - \frac{g_c^I - g_c^K}{\frac{k_c}{2\chi i_c}-1} + \chi a_d^2 - \chi a_c^2}{P_{k_d}\,(1-2\chi a_d)}. \tag{A17}$$

For negative investment, an analogous derivation yields

$$\text{MAC} = -\lambda^S c^\eta = \frac{y_{k_d} + \chi_s a_d^2 - 2\chi_s a_d \left(r + g_L + \eta\frac{\dot{c}}{c} + \delta - g_d^I + g_d^K\right)}{P_{k_d} - P_{k_s}}. \tag{A18}$$

To give insight into the meaning of the shadow price of k_s, we integrate its costate equation,[16]

$$\lambda^{k_s}(t) = \int_t^\infty e^{-(r+\delta+g_L)(\tau-t)} \lambda^S(\tau)\, P_{k_s}(\tau)\, d\tau. \tag{A19}$$

This shows that the shadow price is positive and corresponds to the difference between the highest and lowest emissions intensity of dirty capital, valued at the social cost of carbon, discounted at the sum of the discount rate and the decay rate of capital. λ^{k_s} is initially decreasing because P_{k_s} decreases as we go down the stranding function (while k_s is increasing). Thereafter, λ^{k_s} can be increasing because $\frac{\partial P(\tau)}{\partial k_s(\tau)}$ increases at rate $g + g_L - g_\psi$. The shadow price of k_s has the following steady state,

$$\lambda^{k_s} = \frac{\lambda^S P_{k_s}}{r + \delta + g_L}. \tag{A20}$$

In the stochastic model, we assume that $g_\psi(t) = g_\psi(0) + (g + g_L - g_\psi(0))\left(1 - \exp\left(-g_{g_\psi} t\right)\right)$, where g_{g_ψ} is a parameter governing the convergence speed of g_ψ. This converges to $g + g_L$ to avoid ever-increasing P_{k_s} (it will also stabilise P_{k_d}). This helps convergence of the algorithm as it allows for a well-defined steady state.

In the steady state, the transversality condition of λ^S imposes $k_d = 0$, therefore, $i_d = 0$, and using $\dot{\lambda^{kd}} = 0$ gives the steady state shadow price of dirty capital

$$\lambda^{k_d} = \frac{c^{-\eta} y_{kd} + \lambda^S P_{k_d}}{r + \delta + g_L}, \tag{A21}$$

[16] The general solution to the differential equation includes a term $ke^{(r+\delta+g_L)t}$, with k the integrating constant. From the transversality condition $\lim_{t\to\infty} \lambda^{k_s} e^{-(r-g)t} = 0$, it follows that $k = 0$, which we make explicit in the numerical optimisation by imposing bounds $[0, 1]$ on λ^{k_s}.

which is positive because $y_{k_d} > r+\delta+g_L$. Finally, in the steady state we also have $\lambda^{k_s} = \lambda^{k_d}$.

A.1 Dirty capital, stranding and emissions

This subsection characterises how heterogeneous dirty capital translates into aggregate emissions and how stranding affects marginal emissions.

Dirty capital is heterogeneous in its carbon intensity, indexed by i. The stranding function has the exponential form

$$\psi(\tilde{k}_{d,i}) = e^{-g_\psi t}\left(\psi_1 e^{-\psi_2 \tilde{k}_{d,i}} + \psi_3\right). \tag{A22}$$

Integrating over dirty capital yields emissions. Initial emissions are:

$$P(0) = \frac{\psi_1}{\psi_2}\left(1 - e^{-\psi_2 k_d(0)}\right) + \psi_3 k_d(0). \tag{A23}$$

Dirty capital decays at rate δ.[17] Emissions at time t are hence

$$\begin{aligned} P(t) = A(0)L(0)\exp((g+g_L-g_\psi)t) \\ \times \left\{\frac{\psi_1}{\psi_2}\frac{k_d+k_s}{k_d(0)}\left[\exp\left(-\psi_2 k_d(0)\frac{k_s}{k_d+k_s}\right) - \exp(-\psi_2 k_d(0))\right]\right. \\ \left. + \psi_3 k_d\right\}. \end{aligned} \tag{A24}$$

The marginal effect of dirty capital on emissions is:

$$\begin{aligned} \frac{\partial P}{\partial k_d} = A(0)L(0)\exp((g+g_L-g_\psi)t) \\ \times \left\{\frac{\psi_1}{\psi_2}\exp\left(-\psi_2 k_d(0)\frac{k_s}{k_d+k_s}\right)\left[\frac{1}{k_d(0)} + \psi_2\frac{k_s}{k_d+k_s}\right] - \frac{\psi_1}{\psi_2 k_d(0)}\exp(-\psi_2 k_d(0)) + \psi_3\right\}. \end{aligned}$$

For stranding, the marginal emissions reduction is larger. The reduction from stranding one unit of capital is:

$$\frac{\partial P}{\partial k_d} - \frac{\partial P}{\partial k_s} = A(0)L(0)\exp((g+g_L-g_\psi)t)\left(\psi_1\exp\left(-\psi_2 k_d(0)\frac{k_s}{k_d+k_s}\right) + \psi_3\right). \tag{A25}$$

[17] Due to this decay, the carbon intensity of a given type evolves as $\psi\left(k_{d,i}\frac{k_d(0)}{k_d+k_s}\right)$. During periods of negative investment, $\frac{k_d(0)}{k_d+k_s} = e^{\delta t}$. During periods of positive investment, newly installed capital is of average intensity. As a result, the least carbon-intensive type remains $\psi_1 e^{-\psi_2 k_d(0)} + \psi_3$.

This uses:

$$\begin{aligned}\frac{\partial P}{\partial k_s} = & -A(0)L(0)\exp((g+g_L-g_\psi)t) \\ & \times\left\{-\frac{\psi_1}{\psi_2}\frac{1}{k_d(0)}\left[\exp\left(-\psi_2 k_d(0)\frac{k_s}{k_d+k_s}\right)-\exp(-\psi_2 k_d(0))\right]\right. \\ & \left.+\psi_1\frac{k_d}{k_d+k_s}\exp\left(-\psi_2 k_d(0)\frac{k_s}{k_d+k_s}\right)\right\}.\end{aligned}$$

The term $\frac{\partial P}{\partial k_s}$ declines rapidly over time, despite the positive growth rate $g + g_L - g_\psi$, because the marginal impact of additional stranding decreases as k_s becomes large. In addition, $\frac{k_s}{k_d+k_s} \to 1$. This behaviour is partly driven by the softplus smoothing but is not quantitatively important, as $\frac{\partial P}{\partial k_s}$ only affects λ^{k_s}, which becomes irrelevant after the initial stranding phase.

Appendix B Decentralised competitive equilibrium

This appendix presents a decentralised competitive equilibrium corresponding to the model. All quantities are expressed in levels. This choice is natural for the decentralised equilibrium, which describes ownership, prices, market clearing, and policy instruments measured in physical and monetary units (e.g. tonnes of emissions, carbon prices per tonne).

Markets are competitive and all technologies exhibit constant returns to scale in priced inputs, implying zero static profits. Households choose consumption; firms solve dynamic investment problems with adjustment costs, and the shadow values of installed capital determine investment through standard q-theory conditions; government policy internalises the climate externality and the learning-by-doing externality arising from abatement.

Households

The representative household supplies labour inelastically and owns all productive capital stocks. Its financial wealth W is the market value of final-good, clean-energy, and dirty-energy capital,

$$W = q_f K_f + q_c K_c + q_d K_d, \tag{A26}$$

where q_j denotes the price of one unit of capital of type j. Stranded capital K_s earns no rental income and has no market value, so it does not enter (A26).

The household chooses a consumption path $\{C\}_{t\geq 0}$ to maximise

$$\int_0^\infty e^{-\rho t} \frac{(C/L)^{1-\eta}}{1-\eta} L \, dt, \tag{A27}$$

subject to the wealth-accumulation equation

$$\dot{W} = \hat{r} W + wL + \mathrm{TR} - C, \tag{A28}$$

where w is the wage and TR are lump-sum government transfers. In equilibrium, no arbitrage implies that all assets yield the same return $\hat{r}$, which is the discount rate used by firms. The household's optimality condition for consumption is the standard Euler equation

$$\hat{r} = \rho + \eta g_{C/L}, \tag{A29}$$

where $g_{C/L}$ is the growth rate of consumption per capita.

Final-Good Firms

Final-good firms produce output using

$$Y = (AL)^{1-\alpha}\left(\varphi_f K_f^{\epsilon_f} + \varphi_E E^{\epsilon_f}\right)^{\alpha/\epsilon_f}\Omega, \tag{A30}$$

where climate damages are

$$\Omega = \exp\!\left(-\frac{\gamma}{2}T^2\right), \qquad T = \zeta S, \qquad \dot{S} = P. \tag{A31}$$

Final-good firms choose labour, energy inputs, and investment to maximise the present discounted value of profits. Let q_f denote the shadow value of installed final-good capital. The firm solves

$$V_f = \max_{\{L,E,I_f\}} \int_0^\infty e^{-\int_0^t \hat{r}(s)\,ds}\Big[Y - wL - p_E E - I_f\Big]\,dt, \tag{A32}$$

subject to the capital-accumulation equation

$$\dot{K}_f = I_f - \chi\frac{I_f^2}{K_f} - \delta K_f. \tag{A33}$$

The static FOCs for inputs are

$$w = \frac{\partial Y}{\partial L}, \qquad p_E = \frac{\partial Y}{\partial E}. \tag{A34}$$

The investment FOC is

$$1 = q_f\left(1 - 2\chi\frac{I_f}{K_f}\right), \tag{A35}$$

and the costate equation is

$$\dot{q}_f = \left(\hat{r} + \delta - \chi\frac{I_f^2}{K_f^2}\right)q_f - \frac{\partial Y}{\partial K_f}. \tag{A36}$$

Energy Sector

An integrated energy firm owns and operates both clean and dirty capital. The integrated representation makes transparent how the composition of energy supply responds to prices and policy, because the trade-off between clean and dirty capital is internalised within a single optimisation problem. However, an equivalent equilibrium could be obtained with separate clean- and dirty-energy firms.

Energy output is produced using the CES aggregator

$$E = [\varphi_d(K_d + \xi)^{\epsilon_t} + \varphi_c K_c^{\epsilon_t}]^{1/\epsilon_t}, \tag{A37}$$

where the time-varying substitution parameter ϵ_t is determined by

$$\epsilon_t = 1 - \frac{\mathrm{TC}}{\sigma_0}, \qquad \mathrm{TC} = \omega e^{-g_\sigma t} + (1-\omega)\left(\frac{M}{M_0}\right)^{-\varrho}. \tag{A38}$$

Cumulative experience evolves according to

$$M = M_0 + \int_0^t \Big(P_{\mathrm{BAU}} - P(u)\Big)\, du. \tag{A39}$$

Emissions are generated by dirty capital and depend on the composition of the remaining dirty-capital stock. Let K_s denote cumulative stranded dirty capital, interpreted as the amount of the highest-emissions-intensity dirty capital that has been retired. Then aggregate emissions can be written in reduced form as

$$P = P(K_d, K_s), \tag{A40}$$

with

$$P_{K_d} \equiv \frac{\partial P}{\partial K_d} > 0, \qquad P_{K_s} \equiv \frac{\partial P}{\partial K_s} < 0. \tag{A41}$$

The dependence on K_s captures the idea that stranding the dirtiest units of dirty capital lowers the emissions intensity of the remaining stock. This reduced-form representation is the decentralised counterpart of the planner's emissions block.

The energy firm maximises the present discounted value of profits,

$$V_E = \max_{\{I_c, I_d\}} \int_0^\infty e^{-\int_0^t \hat{r}(s)\, ds} \Big[p_E E - \tau_P P + s_A\Big(P_{\mathrm{BAU}} - P\Big) - I_c - \mathbf{1}_{(I_d>0)} I_d - \Theta\Big]\, dt, \tag{A42}$$

subject to

$$\dot{K}_c = I_c - \chi \frac{I_c^2}{K_c} - \delta K_c, \tag{A43}$$

$$\dot{K}_d = I_d - \mathbf{1}_{(I_d>0)}\, \chi \frac{I_d^2}{K_d} - \delta K_d, \tag{A44}$$

$$\dot{K}_s = \max\{-I_d, 0\} - \delta K_s, \tag{A45}$$

where τ_P is a carbon tax, s_A is an abatement subsidy (more on these just below), and $\Theta = \mathbf{1}_{(I_d<0)} \chi_s I_d^2 / K_d$ is the stranding cost.

Let q_c, q_d, and q_s denote the shadow values of clean, dirty, and stranded capital. Because negative dirty investment simultaneously reduces active dirty capital and increases stranded capital, the marginal condition for $I_d < 0$ depends on both q_d and q_s. The investment FOCs are

$$1 = q_c \left(1 - 2\chi \frac{I_c}{K_c}\right), \tag{A46}$$

$$1 = q_d \left(1 - 2\chi \frac{I_d}{K_d}\right), \qquad I_d > 0, \tag{A47}$$

$$0 = q_d - q_s - 2\chi_s \frac{I_d}{K_d}, \qquad I_d < 0, \tag{A48}$$

and the costate equations are

$$\dot{q}_c = \left(\hat{r} + \delta - \chi \frac{Ic^2}{Kc^2}\right) q_c - p_E \frac{\partial E}{\partial K_c}, \tag{A49}$$

$$\dot{q}_d = \left(\hat{r} + \delta - \mathbf{1}_{Id>0}\chi \frac{Id^2}{Kd^2}\right) q_d - \mathbf{1}_{I_d<0}\chi_s \frac{I_d^2}{K_d^2} - p_E \frac{\partial E}{\partial K_d} + (\tau_P - s_A)\frac{\partial P}{\partial K_d}, \tag{A50}$$

$$\dot{q}_s = (\hat{r} + \delta) q_s + (\tau_P - s_A)\frac{\partial P}{\partial K_s}. \tag{A51}$$

Government

The government levies the tax τ_P on emissions P and pays the subsidy s_A per unit of abatement

$$A = P_{\text{BAU}}(0) - P.$$

Its flow budget constraint is

$$\tau_P P = s_A A + \text{TR}, \tag{A52}$$

with transfers rebated lump-sum to households.

Induced technical change operates through cumulative abatement, via its effect on TC and hence on ϵ_t. We assume it is fully externalised. The planner's shadow value of emissions therefore contains two components: marginal climate damages and the marginal value of induced innovation. In the decentralised economy we represent these using a "damage-only" carbon tax τ_P and an abatement subsidy s_A paid per unit of emissions reduced relative to a BAU benchmark. An equivalent representation is a single adjusted carbon price equal to the full marginal social value of emissions; both formulations implement the same wedge $(\tau_P + s_A)$ in the dirty-capital condition.

Market Clearing and Feasibility

Capital and labour markets clear. Goods market clearing is given by

$$C = Y - I_f - I_c - \mathbf{1}_{(I_d>0)} I_d - \Theta. \tag{A53}$$

When $I_d < 0$, dirty capital is scrapped without salvage value and disinvestment does not add resources for consumption or investment.

Competitive Equilibrium

Definition 1 (Competitive Equilibrium)**.** *Given policy paths $\{\tau_P, s_A\}$ and initial states $K_f(0)$, $K_c(0)$, $K_d(0)$, $K_s(0)$, $S(0)$, a competitive equilibrium consists of allocation paths, price paths, and transfers such that:*

1. *households choose consumption to maximise lifetime utility subject to their wealth-accumulation equation and take prices and transfers as given;*
2. *final-good firms choose inputs and investment to maximise their value, taking prices as given and satisfying their FOCs and costate equation for capital accumulation;*
3. *the energy firm chooses clean and dirty investment to maximise its value, taking prices and policy instruments as given and satisfying its investment FOCs and costate equations, including adjustment and stranding costs;*
4. *the government budget constraint holds;*
5. *labour, capital, and goods markets clear; and*
6. *state variables evolve according to the laws of motion specified above.*

Decentralisation

Proposition 1 (Decentralisation of the Planner's Allocation)**.** *Let $(\cdot)^\star$ denote an allocation that solves the planner's problem in the main text. Let $u'(c^\star) = c^{\star -\eta}$ denote the marginal utility of consumption along the planner's allocation, and let λ_S be the planner's shadow value of cumulative emissions. Decompose*

$$\lambda_S = \lambda_S^{\mathrm{SCC}} + \lambda_S^{\mathrm{TC}},$$

where λ_S^{SCC} captures marginal climate damages and λ_S^{TC} captures the marginal value of induced technical change. Define policy instruments

$$\tau_P^{\text{damage}} = -\frac{\lambda_S^{\text{SCC}}}{c^{\star-\eta}}, \qquad s_A^{\star} = -\frac{\lambda_S^{\text{TC}}}{c^{\star-\eta}}. \tag{A54}$$

Then the planner's allocation can be decentralised as a competitive equilibrium under the policy paths $\left(\tau_P^{\text{damage}}, s_A^{\star}\right)$.

Equivalently, the same allocation can be decentralised using a single adjusted carbon price equal to the full marginal social value of emissions.

Proof sketch. To compare the decentralised equilibrium to the social planner's solution, convert the shadow prices q_j in consumption units to shadow prices λ_j expressed in utils,

$$q_j = \lambda^j c^{\eta} = \lambda^j \left(\frac{(C/L)}{A}\right)^{\eta} \tag{A55}$$

$$\Leftrightarrow \frac{\dot{q_j}}{q_j} = \frac{\dot{\lambda^j}}{\lambda^j} + \eta\frac{\dot{c}}{c}, \tag{A56}$$

where we use $g_{C/L} = \frac{\dot{c}}{c} + g$. With this transformation and using $\hat{r} = r + g_L + \eta\frac{\dot{c}}{c}$, the firms' FOCs (A35, A46, A47 and A48) are identical to the planner's (A3 A4, A6). Similarly, costate equations A36, A49,A50, and A51 result in costate equations A7, A8, A9 and A10. □

The decentralisation result assumes that private agents and the social planner share the same discount rate. If private discounting differed from the social rate, additional intertemporal wedges would arise and the carbon-price path would no longer decentralise the planner's allocation (Barrage, 2018). The result also assumes that policy instruments are credible and that the policy-maker can commit to future carbon taxes and abatement subsidies. Because preferences are time-separable and discounting is exponential, the planner's problem itself is dynamically consistent (Iverson and Karp, 2021). However, as in standard investment models with irreversible capital and adjustment costs, lack of policy commitment could distort private investment incentives through hold-up effects or induce deviations from the Ramsey carbon price path. We abstract from these issues here.

Appendix C Further details on calibration

C.1 Calibration table

Table A1: Parameters and initial values in the central scenario

Parameter	Symbol	Value	Source
Production			
Initial TFP value	A_0	8.09	Calibrated to match Y_0
TFP long-run growth rate	g_A	0.02/yr	By assumption
Initial population	L_0	8.16 bn	UN (2024)
Population growth rate	g_L	0.0029/yr	UN (2024)
Output elasticity of capital-energy	α	0.3	Standard
Share of E in final good production	φ_E	0.00005603	Implied by $y_{kf} = y_{kc}$
Share of K_f in final good production	φ_f	0.999944	Implied by $\varphi_f = 1 - \varphi_E$
Elasticity of K_f-E substitution	σ_f	0.2	In line with literature
Energy and emissions			
Share of K_d in energy production	φ_d	0.696	Implied by $y_{kc} = y_{kd}$
Share of K_c in energy production	φ_c	0.304	Implied by $\varphi_d = 1 - \varphi_c$
Initial elasticity of K_d-K_c substitution	σ_0	2	Papageorgiou et al. (2017)
Initial GHG emissions	P_0	53.2 GtCO2e	Crippa et al. (2025)
Emission intensity premium	ψ_1	6.120	Estimated (Section 4)
Emission intensity decay rate	ψ_2	0.729	Estimated (Section 4)
Emission intensity floor	ψ_3	0.767	Estimated (Section 4)
Emission intensity exogenous decline rate	g_ψ	0.01/yr	World Bank (2026)
Capital			
Initial final good capital	$K_{f,0}$	\$321.3 tn	Implied by balanced growth path
Initial dirty capital	$K_{d,0}$	\$58.4 tn	Estimated (Section 4)
Dirty capital shifter	ξ	0.01	By assumption
Initial clean capital	$K_{c,0}$	\$11.3 tn	Implied by $K_{d,0}$ and IEA (2025)
Asset lifetime	LT	34.65 years	Calibrated to match δ_d
Depreciation rate	δ	0.04	In line with literature
Adjustment cost parameter	χ	1	Calibrated (Section 4)
Stranding cost parameter	χ_s	5	Calibrated (Section 4)
Std. dev. of capital Brownian motion	σ_K	$0.015/\text{yr}^{1/2}$	van den Bremer and van der Ploeg (2021)
Clean technological progress			
Rate of exogenous technological progress	g_σ	0.027/yr	Coppens et al. (2025)
Exogenous learning weight parameter	ω	0.5	By assumption
Cumulative abatement at t_0	M_0	50 GtCO_2e	Coppens et al. (2025)
Elasticity of TC to cumulative abatement	ϱ	0.211	Coppens et al. (2025)
Damages and warming			
Temperature increase in 2024	T_0	1.38 °C	WMO (2025)
TCRE	ζ	$0.0006\,°\text{C}/\text{GtCO}_2$	IPCC (2021)
Climate damage function parameter	γ	0.0201	Howard and Sterner (2017)
Std. dev. of T Brownian motion	σ_T	$0.0265\,°\text{C}/\text{yr}^{1/2}$	IPCC (2021)
Preferences			
Elasticity of marginal utility of consumption	η	1.35	Drupp et al. (2018)
Relative risk aversion	RRA	4	Barro (2009)
Utility discount rate	ρ	0.011/yr	Drupp et al. (2018)

C.2 Construction of firm-level stranding metrics

Selection of sectors to analyse

We identify the most polluting sectors using Eurostat data on EU sectoral GHG emissions, the only major source reporting emissions using NACE categories, and retain productive sectors emitting more than 50 $MtCO_2e$, excluding households. This yields the following 22 NACE level 2 sectors: B05–09 Mining and quarrying; C10–12 Manufacture of food, beverage and tobacco; C19 Manufacture of coke and refined petroleum products; C20 Manufacture of chemicals and chemical products; C23 Manufacture of other non-metallic mineral products (including cement); C24 Manufacture of basic metals (including iron, steel and aluminium); D35 Electricity, gas, steam and air conditioning supply; E37–39 Waste management; F41–43 Construction; H49 Land transport and pipelines; H50 Water transport; H51 Air transport. We exclude agriculture because of concerns about the comprehensiveness of company-level non-CO2 emissions reporting in Orbis. We validate this sector selection against global sectoral emissions data from Crippa et al. (2025), converted to NACE categories.

Sector-specific scaling factors

This section describes the construction of sector-specific scaling factors that reconcile firm-level emissions from Orbis with sectoral totals from EDGAR and UNFCCC. For each NACE sector j, the goal is to obtain a multiplicative factor s_j such that scaled Orbis emissions match the corresponding external benchmark, under the assumption that the share of emissions captured by Orbis is representative of the uncovered remainder.

We start from the EDGAR emissions dataset (Crippa et al., 2025), which reports 2024 country-level and global GHG emissions by sector using IPCC (2006) categories. The emissions data reported by UNFCCC (2026) are less recent and less geographically comprehensive, but in some cases more disaggregated across IPCC (2006) categories. Whenever available, we therefore use UNFCCC sub-sector *shares* to further split broad EDGAR sector aggregates. Before doing so, we adjust for differences in the treatment of agriculture and LULUCF emissions between the EDGAR/IPCC accounting framework and the UNFCCC Common Reporting Framework (CRF).

Next, we map the resulting IPCC/CRF sectoral emissions to the NACE Rev. 2 level-2 categories used in Orbis, using the CRF–NACE correspondence table in Eurostat (2015, Annex I). Several adjustments are needed for consistency. First, when a CRF category maps to multiple NACE codes, we split its emissions equally across them. Second, the table includes entries mapped to the label 'CS', denoting country-specific reporting categories

Table A2: Sector-specific scaling factors s_j (multipliers) applied to Orbis firm emissions. Shown values are rounded to the third decimal place.

NACE sector	s_j	NACE sector	s_j	NACE sector	s_j
B5–10	3.244	C23	2.304	F41–43	0.470
C10–12	2.937	C24	1.006	H49	10.716
C19	1.477	D35	5.595	H50	2.271
C20	2.462	E37–39	13.125	H51	2.380

without a harmonised CRF definition. We conservatively exclude 'CS' codes when splitting emissions. Whenever a CRF subcategory is associated *uniquely* with a 'CS' label, we exclude the corresponding emissions altogether; this implies a loss of about 2% of total emissions. Finally, the label '01–99' also appears, for which insufficient information for sector-level allocation is available. We therefore treat it in the same way as 'CS'. In practice, this causes no additional loss of emissions, because '01–99' always appears alongside other well-defined sectoral codes.

Let E_j^{EDGAR} denote the resulting benchmark emissions allocated to NACE sector j, and let $E_j^{\mathrm{ORBIS}} = \sum_{i \in j} e_i$ the corresponding unscaled ORBIS sector total. We define the sectoral scaling factor as

$$s_j \equiv \frac{E_j^{\mathrm{EDGAR}}}{E_j^{\mathrm{ORBIS}}}. \tag{A57}$$

Scaled firm emissions are then given by $e_i^{\mathrm{scaled}} = s_{j(i)}\, e_i$, where $j(i)$ is firm i's NACE sector. By construction, $\sum_{i \in j} e_i^{\mathrm{scaled}} = E_j^{\mathrm{EDGAR}}$ for each sector j. Table A2 reports the resulting scaling factors.

Asset age structure

For some companies, physical asset values are observed only for either K_n or K_g. In these cases, we impute the missing value so that implied asset age equals the mean age among firms for which both are observed, separately by asset sub-category. In total, 618 companies have at least one imputed capital value, mostly in 'Other plant, property and equipment'.

LT and δ must be determined jointly to reconcile: (i) the linear depreciation structure applied to Orbis data; and (ii) the exponential depreciation structure in the model. We do so by matching the asset economic half-life across the two settings, i.e. the number of years required for the asset's value to fall to half its original level. Since economic half-life is $LT/2$ under linear depreciation and $\ln(2)/\delta$ under exponential depreciation, we obtain $LT = 2\ln(2)/\delta$, and therefore set $\delta_d = 0.04$ and $LT = 34.65$ years jointly. These values are consistent with both depreciation rates used in models with long-lived physical assets and

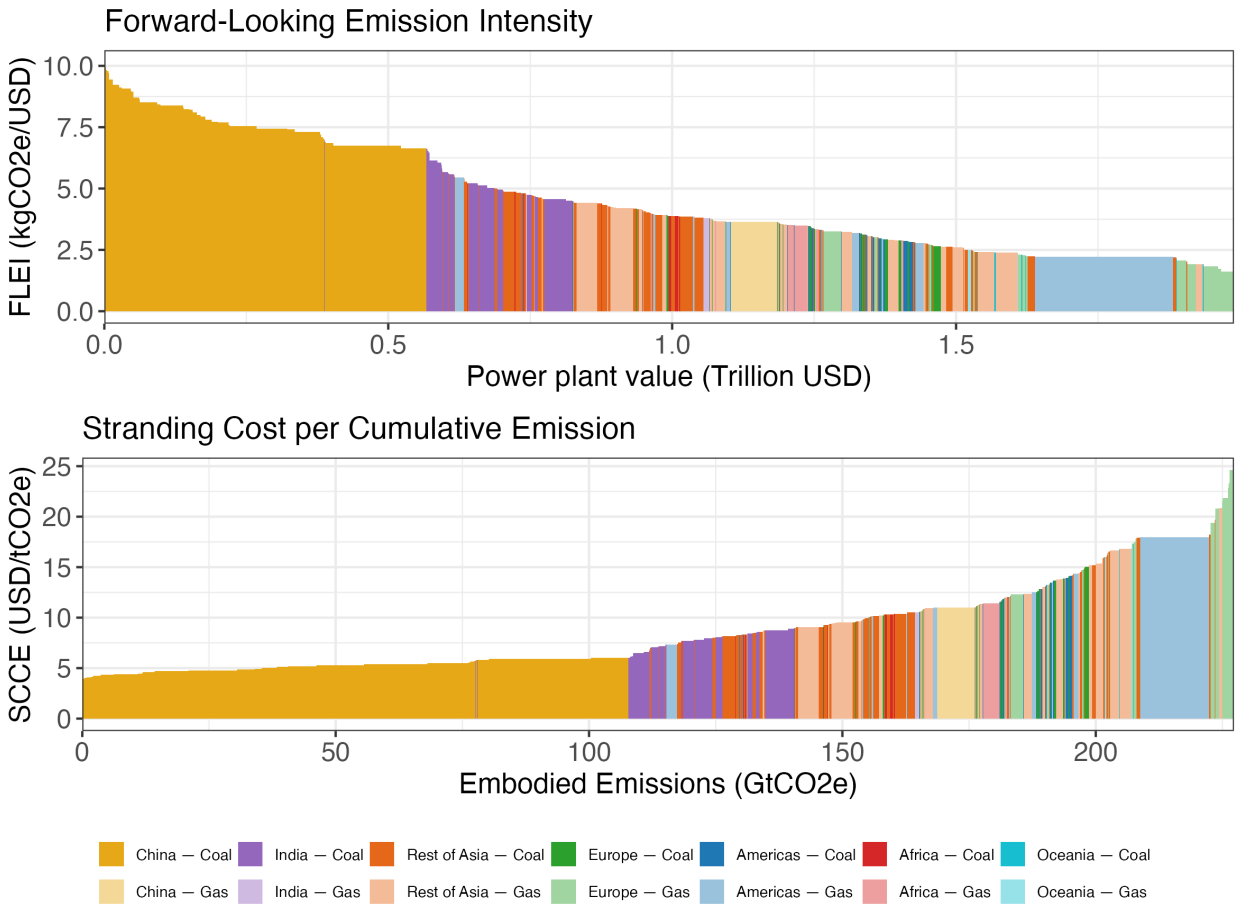

Figure A1: Plant-level forward-looking emissions intensity (FLEI - upper panel) and stranding cost per cumulative emission (SCCE - lower panel). Total observations: 5,073 (coal); 6,951 (gas). The width of each bar is proportional to the monetary value of power plants (upper panel) or the amount of embodied emissions (lower panel). FLEI/SCCE values are winsorised at 99%.

empirical estimates of typical physical capital lifetimes (see for instance IEA, 2020).

C.3 Plant-level analysis

This section replicates the estimation procedure presented in Section 4 using plant-level rather than company-level data. A fully comprehensive plant-level dataset encompassing all productive sectors would be ideal, but plant-level information remains sparse outside power generation. We therefore provide complementary evidence in support of our company-level results using the Global Coal Plant Tracker and the Global Oil and Gas Plant Tracker of GEM (2026). These dataset report generation capacity and commissioning year for more than 5,073 operating coal plants and 6,951 gas plants worldwide. Given our chosen value

of LT, the commissioning year allows us to infer each plant's remaining lifetime, $LT - a$. While GEM (2026) also offers asset-level estimated emissions for coal plants, we estimate emissions for gas plants ourselves, combining capacity data with country- or region-specific utilisation factors from IPCC (2014) and a single technology-wide lifecycle emission factor of 490 gCO_2eq/kWh (IPCC, 2006).

To approximate plant monetary values, we convert capacity into capital values using Overnight Capital Costs (OCC) estimates. OCCs cover pre-construction, construction, and contingency costs, and are routinely used in technology comparisons and levelised cost of electricity (LCOE) calculations. Using region- and technology-specific OCC estimates from von Dulong (2023), we construct $K_n = GW \times OCC \times (LT - a)/LT$.[18]

We then follow the same steps as in Section 4. Results are shown in Figure A1, where colour indicates plant location and colour transparency the technology.. Despite the different calculation strategy, plant-level results closely match the company-level patterns discussed in the main text. Coal and gas power plants represent the very left end of both the company-level SCCE and FLEI plots, where companies owning fossil power plants would appear. We confirm that abating emissions through stranding appears to be relatively inexpensive: almost the entire global fleet of coal and gas plants, when considering current and future emissions, can be stranded at a cost below 20$/$tCO_2e$. The figure also shows that most fossil power plants – including those that should be stranded first – are located in China, India and other Asian countries, raising geopolitical questions about the feasibility of a global dirty-asset stranding strategy.

C.4 Calibration on the BAU steady state

We use steady-state conditions to pin down the two CES share parameters, φ_d and φ_E, as well as the initial final-good capital stock $k_f(0)$ (and hence $K_f(0)$).

Expressed in units of effective labour, the equations of motion for the three productive capital stocks can be written as

$$\dot{k}_j = i_j - \chi \frac{i_j^2}{k_j} - (\delta + g + g_L)k_j.$$

We also define the ratio of dirty to final-good capital as $\theta_d \equiv (k_d + \xi)/k_f$ and the ratio of

[18] von Dulong (2023) reports overnight capital costs for coal (first value reported) and gas/oil power plants (second value): 3,095/812$/kW (Australia); –/906$/kW (Canada); 2,000/1,000$/kW (EU); 2,419/1,109$/kW (Japan); 1,151/973$/kW (Korea); –/601$/kW (Mexico); 2,100/1,000$/kW (US); 2,153/914$/kW (Other OECD countries); 2,189/847$/kW (Brazil); 800/560$/kW (China); 1,200/700$/kW (India); 1,396/702$/kW (Other non-OECD countries). OCC values only for gas/oil plants are available for two countries: 906$/kW (Canada); 601$/kW (Mexico)

clean to final-good capital as $\theta_c \equiv k_c/k_f$.

In the steady state of the model expressed in effective-labour units – equivalently, on the balanced growth path in levels – all productive capital stocks satisfy the same law of motion. Hence, the investment-capital ratio is common across sectors, so that $a_j \equiv i_j/k_j = a$ for all $j \in \{f, c, d\}$. The steady-state condition $\dot{k}_j = 0$ then implies

$$a = \chi a^2 + (\delta + g + g_L) \Leftrightarrow a = \frac{1 - \sqrt{1 - 4\chi\,(\delta + g + g_L)}}{2\chi}. \tag{A58}$$

Combining this with the investment FOCs and the costate equations for productive capital, and evaluating the dirty-capital condition in the non-stranding BAU regime, yields the common steady-state return condition

$$y_{kd} = y_{kc} = y_{kf} = \frac{\delta + r + g_L - \chi a^2}{1 - 2\chi a} \equiv R^*. \tag{A59}$$

Since $\varphi_c = 1 - \varphi_d$, the condition $y_{kc} = y_{kd}$ yields

$$\varphi_d = \left[\left(\frac{\theta_c}{\theta_d}\right)^{1-\epsilon} + 1\right]^{-1}. \tag{A60}$$

Rewriting production as

$$y = k_f^{\alpha}\left[\varphi_f + \varphi_E\left(\varphi_c\theta_c^{\epsilon} + \varphi_d\theta_d^{\epsilon}\right)^{\epsilon_f/\epsilon}\right]^{\alpha/\epsilon_f}\Omega, \tag{A61}$$

and using $\varphi_f = 1 - \varphi_E$, the condition $y_{kf} = y_{kc}$ yields the energy share parameter

$$\varphi_E = \frac{1}{1 + \left(\varphi_c\theta_c^{\epsilon} + \varphi_d\theta_d^{\epsilon}\right)^{\frac{\epsilon_f}{\epsilon}-1}\varphi_c\theta_c^{\epsilon-1}}. \tag{A62}$$

Finally, the condition $y_{kf} = R^*$ pins down the steady-state value of final-good capital:

$$k_f = \left[\frac{\alpha\varphi_f\left(\varphi_f + \varphi_E\left(\varphi_c\theta_c^{\epsilon} + \varphi_d\theta_d^{\epsilon}\right)^{\epsilon_f/\epsilon}\right)^{\alpha/\epsilon_f - 1}\Omega}{R^*}\right]^{\frac{1}{1-\alpha}}. \tag{A63}$$

This shows that damages also have a decreasing effect on the steady state capital, via a lower incentive on investment. Combined with equation A61, this shows that the damage factor in the steady state production is raised to a power $\frac{1}{1-\alpha}$. For example, for $\alpha = 0.3$ an initial damage of 3% becomes 4.3% after taking into account the feedback on capital.

Appendix D HJB equation, Itô's lemma and optimality conditions

Consumption per unit of effective labour is a function of six state variables $\{k_f, k_d, k_c, k_s, S, t\}$ and three control variables $\{i_f, i_d, i_c\}$:

$$c(k_f, k_d, k_c, k_s, S, t, i_f, i_d, i_c) = y - i_f - i_c - \mathbf{1}_{i_d>0} i_d - \mathbf{1}_{i_d<0} \chi_s \frac{i_d^2}{k_d}. \tag{A64}$$

This defines the value function in Equation (20) and the recursive aggregator in Equation (21).

The HJB equation is

$$\max_{i_f, i_d, i_c} \left\{ f + \frac{1}{dt} \mathbb{E}[dV(k_f, k_d, k_c, k_s, S, t)] \right\} = 0,$$

where discounting is already incorporated through the transformation to effective labour units.

Applying Itô's lemma yields

$$\begin{aligned}
0 \;=\; & \max_{i_f, i_d, i_c} \Big\{ f \\
& + V_{k_f} \left(i_f - \chi \frac{i_f^2}{k_f} - (\delta + g_L + g) k_f \right) \\
& + V_{k_d} \left(i_d - \mathbf{1}_{i_d>0} \chi \frac{i_d^2}{k_d} - (\delta + g_L + g) k_d \right) \\
& + V_{k_c} \left(i_c - \chi \frac{i_c^2}{k_c} - (\delta + g_L + g) k_c \right) \\
& + V_{k_s} \left(\mathbf{1}_{i_d<0} (-i_d) - (\delta + g_L + g) k_s \right) \\
& + \frac{1}{2} V_{k_d k_d} k_d^2 \sigma_{K_d}^2 + \frac{1}{2} V_{k_c k_c} k_c^2 \sigma_{K_c}^2 + \frac{1}{2} V_{k_f k_f} k_f^2 \sigma_{K_f}^2 \\
& + V_S P + \frac{1}{2} V_{SS} S^2 \sigma_T^2 \\
& + V_t \Big\}.
\end{aligned}$$

Optimality conditions. The first-order condition with respect to dirty investment, away from the kink at $i_d = 0$, is

$$f_c \, c_{i_d} + V_{k_d} \left(1 - \mathbf{1}_{i_d>0} 2\chi a_d \right) - V_{k_s} \, \mathbf{1}_{i_d<0} = 0. \tag{A65}$$

From Equation (A64),

$$c_{i_d} = -\mathbf{1}_{i_d>0} - \mathbf{1}_{i_d<0} 2\chi_s a_d. \tag{A66}$$

Combining these expressions characterises the optimality condition for dirty investment. The corresponding conditions for clean and final capital are analogous but do not involve stranding terms.

The derivative of the recursive aggregator with respect to consumption is

$$f_c = \hat{\rho} c^{-\eta} \left[(1 - \mathrm{RRA})V\right]^{\frac{\eta - \mathrm{RRA}}{1 - \mathrm{RRA}}}. \tag{A67}$$

Carbon price. The shadow value of cumulative emissions in utility units is $-V_S$. We denote partial derivatives of the value function by subscripts, e.g. $V_{k_f} = \partial V / \partial k_f$, and higher-order derivatives accordingly, e.g. $V_{Sk_f} = \partial^2 V / (\partial S \partial k_f)$. To obtain the optimal carbon price in consumption units, we differentiate the HJB equation with respect to S and apply the envelope theorem.

Differentiating the HJB with respect to S yields contributions from: (i) the direct effect of S on the aggregator; (ii) the drift terms for the capital stocks; (iii) the diffusion terms; and (iv) the explicit time dependence of the value function. Collecting these terms gives the shadow value of emissions in utility units.

To convert this shadow value into consumption units, we use the marginal utility of consumption together with the first-order condition for clean capital,

$$f_c = V_{k_c} \left(1 - 2\chi a_c\right).$$

Using Equation (A67) to substitute for f_c then yields the normalisation factor Φ in Equation (30). This delivers the stochastic carbon-price expression reported in the main text.

Appendix E Neural nets solution method

Our algorithm is based on Azinovic et al. (2022). Call $x(k_f, k_c, k_d, k_s, T, \tau)$ the vector of state variables, where τ is synthetic time defined as

$$\tau = 1 - e^{-g_\tau t} \Leftrightarrow t = \frac{-\ln(1-\tau)}{g_\tau} \Leftrightarrow \frac{\partial t}{\partial \tau} = \frac{1}{g_\tau(1-\tau)}. \tag{A68}$$

We approximate the policy and value functions using a neural network

$$\mathcal{N}(x) = (i_f, i_c, i_d, \lambda^{k_s}, \lambda^S, v),$$

which maps states into optimal controls, shadow prices, and the value function.

Recursive problem. The recursive problem is

$$v(x_t) = \max_{\{i_f, i_c, i_d\}} \left\{ \frac{c^{1-\eta}}{1-\eta}\Delta t + \frac{e^{-(r-g)\Delta t}}{1-\eta} \left\{ \mathbb{E}\left[\left((1-\eta)\, v\left(x_{t+\Delta t}\right) \right)^{\frac{1-RRA}{1-\eta}} \right] \right\}^{\frac{1-\eta}{1-RRA}} \right\}. \tag{A69}$$

Instead of imposing $x_{t+\Delta t} = g(x_t, i_f, i_c, i_d)$ directly, we treat $x_{t+\Delta t}$ as additional choice variables and impose the laws of motion as constraints. The resulting Lagrangian is

$$\begin{aligned} \mathcal{L} = \max_{i_f, i_c, i_d, x_{t+\Delta t}} \Bigg\{ & \frac{c_t^{1-\eta}}{1-\eta}\Delta t - e^{-(r-g)\Delta t} \left[\mathbb{E}\left[\left(-v\left(x_{t+\Delta t}\right) \right)^{\frac{1-RRA}{1-\eta}} \right] \right]^{\frac{1-\eta}{1-RRA}} \\ & + \sum_t \sum_{i \in \{k_f, k_c, k_d, k_s, S\}} \lambda_t^i \left[Constraint^i \right] \Bigg\}. \end{aligned}$$

First-order conditions. The first-order conditions are[19]

[19] The derivative of the softplus function $sp_\varepsilon(i_d)$ is the sigmoid function $\varsigma(\varepsilon i_d) = \frac{1}{1+\exp(-\varepsilon i_d)}$.

$$
\begin{aligned}
\frac{\partial \mathcal{L}_t}{\partial i_f} &= 0 = \quad c^{-\eta}\left(-1\right) + \lambda^{k_f}\left(1 - 2\chi a_f\right), \\
\frac{\partial \mathcal{L}_t}{\partial i_c} &= 0 = \quad c^{-\eta}\left(-1\right) + \lambda^{k_c}\left(1 - 2\chi a_c\right), \\
\frac{\partial \mathcal{L}_t}{\partial i_d} &= 0 = \quad c^{-\eta}\left(-\varsigma(\varepsilon_1 i_d) - \mathbf{1}_{i_d<0} 2\chi_s a_d\right) \\
&\qquad\qquad + \lambda^{k_d}\left(1 - \mathbf{1}_{i_d>0} 2\chi a_d\right) + \lambda^{k_s}\varsigma\left((-i_d + \varepsilon_2)\varepsilon_1\right)(-1), \\
\frac{\partial \mathcal{L}_t}{\partial k_f(t+\Delta t)} &= 0 = \quad e^{-(r-g)\Delta t}\mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{1-RRA}{1-\eta}}\right]^{\frac{RRA-\eta}{1-RRA}} \mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{\eta-RRA}{1-\eta}} v_{k_f}\left(x_{t+1}\right)\right] - \lambda_t^{k_f}, \\
\frac{\partial \mathcal{L}_t}{\partial k_d(t+\Delta t)} &= 0 = \quad e^{-(r-g)\Delta t}\mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{1-RRA}{1-\eta}}\right]^{\frac{RRA-\eta}{1-RRA}} \mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{\eta-RRA}{1-\eta}} v_{k_d}\left(x_{t+1}\right)\right] - \lambda_t^{k_d}, \\
\frac{\partial \mathcal{L}_t}{\partial k_c(t+\Delta t)} &= 0 = \quad e^{-(r-g)\Delta t}\mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{1-RRA}{1-\eta}}\right]^{\frac{RRA-\eta}{1-RRA}} \mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{\eta-RRA}{1-\eta}} v_{k_c}\left(x_{t+1}\right)\right] - \lambda_t^{k_c}, \\
\frac{\partial \mathcal{L}_t}{\partial k_s(t+\Delta t)} &= 0 = \quad e^{-(r-g)\Delta t}\mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{1-RRA}{1-\eta}}\right]^{\frac{RRA-\eta}{1-RRA}} \mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{\eta-RRA}{1-\eta}} v_{k_s}\left(x_{t+1}\right)\right] - \lambda_t^{k_s}, \\
\frac{\partial \mathcal{L}_t}{\partial S(t+\Delta t)} &= 0 = \quad e^{-(r-g)\Delta t}\mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{1-RRA}{1-\eta}}\right]^{\frac{RRA-\eta}{1-RRA}} \mathbb{E}\left[\left(-v\left(x_{t+1}\right)\right)^{\frac{\eta-RRA}{1-\eta}} v_S\left(x_{t+1}\right)\right] - \lambda_t^{S}.
\end{aligned}
$$

The loss function is constructed as the sum of squared residuals of these first-order conditions together with the Bellman equation.

Envelope conditions. To compute derivatives of the value function, we apply the envelope theorem to the Lagrangian:

$$
\begin{aligned}
v_{k_f} &= \frac{\partial \mathcal{L}_t}{\partial k_f}, \\
v_{k_c} &= \frac{\partial \mathcal{L}_t}{\partial k_c}, \\
v_{k_d} &= \frac{\partial \mathcal{L}_t}{\partial k_d}, \\
v_{k_s} &= \frac{\partial \mathcal{L}_t}{\partial k_s}, \\
v_S &= \frac{\partial \mathcal{L}_t}{\partial S}.
\end{aligned}
$$

These expressions yield the analytical derivatives used in the loss function.

Implementation details. To stabilise training, we rescale variables so that inputs and outputs of the neural network are of order one. In particular, we scale S, λ^S, i_d, i_c and v

by factors ζ, $1/\zeta$, 100, 10 and 0.01, respectively. Using T instead of S as a state variable is equivalent to this rescaling.

We augment the state vector with auxiliary inputs y_{k_f}, y_{k_c}, y_{k_d} and y_S.

To avoid explosive adjustment costs when capital is close to zero, adjustment costs are capped at 70% of investment. This constraint is not binding along the optimal path given the calibration of χ.

We use the first-order conditions to substitute out λ^{k_f}, λ^{k_c} and λ^{k_d}. For example,

$$\lambda^{k_f} = c^{-\eta} \left(1 - 2\chi a_f\right)^{-1},$$

and similarly for k_c and k_d.

We initialise the neural network with random weights. A batch of initial states is drawn around the initial condition and used to generate simulated trajectories. Parameters are updated using stochastic gradient descent. The neural network has three layers with 200, 200 and 100 nodes, GELU activation functions, and dropout of 1%. We use the AdamW optimiser with a learning-rate schedule decreasing from 7×10^{-4} to 10^{-6} over 200,000 episodes. We start with a model where $\eta = RRA = 1.35$ until the model has found a good value for the value function v (after 100,000 episodes). Next, we gradually increase RRA and increase the relative weight of the Bellman equation.

Appendix F Further results

F.1 Impulse response functions

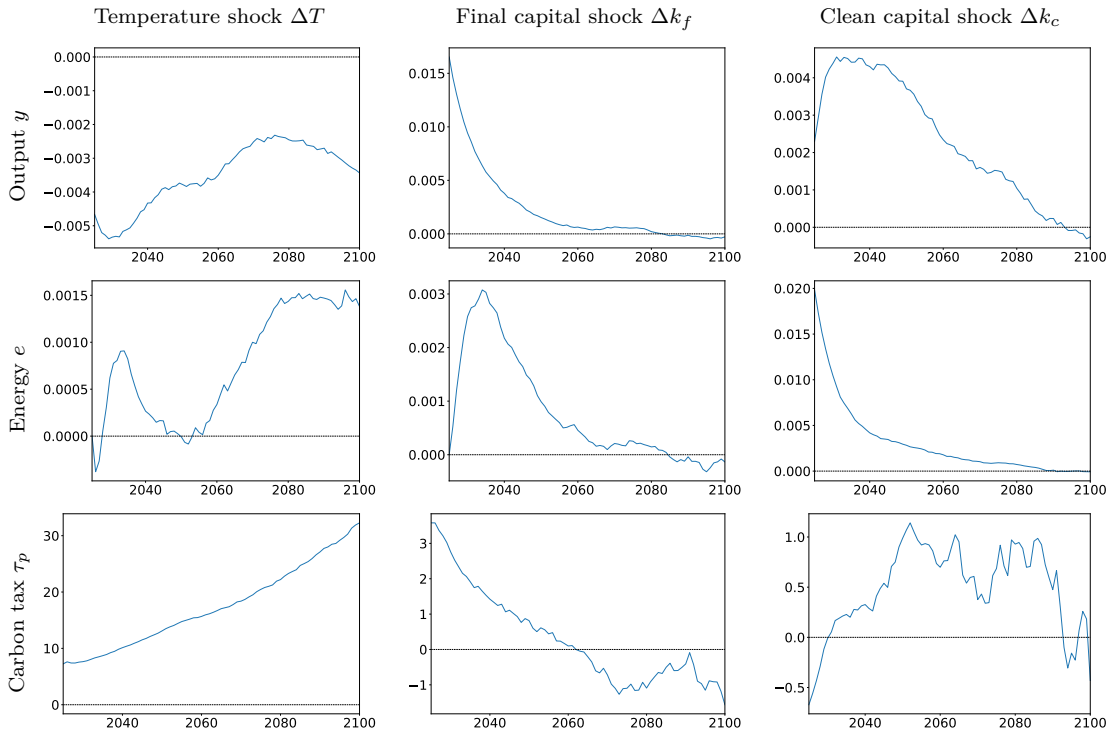

Figure A2: Impulse response functions. Columns correspond to a positive shock on temperature, final capital and clean capital respectively. Rows show the response of output, energy and the carbon tax. Investment is expressed per unit of effective labour. The size of the shocks corresponds to three standard deviations. The capital shocks correspond to 10% of initial capital. The response is the mean difference compared to the baseline for 10,000 Monte Carlo runs.

F.2 Marginal abatement costs

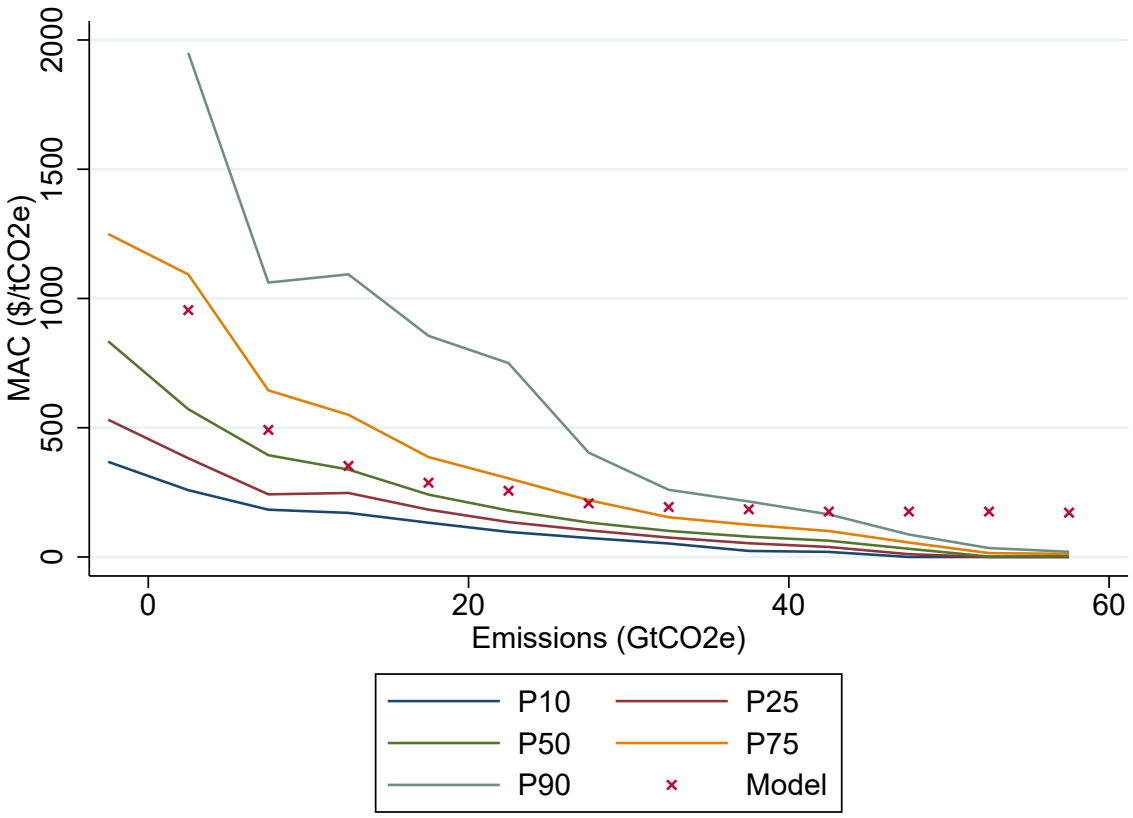

Figure A3: Validation of marginal abatement costs. Markers represent the MACs emerging from the transition in our model. Solid lines represent the median, 25-75th percentiles and 10-90th percentiles of MACs originating from the dataset of IPCC and NGFS models running RCP2.6 scenarios.