Grantham Research Institute on Climate Change and the Environment

Optimal climate policy under exogenous and endogenous technical change: making sense of the different approaches

Leo Coppens, Simon Dietz and Frank Venmans

March 2025 (update)

Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 408

ISSN 2515-5717 (Online)

Grantham Foundation for the Protection of the Environment

The London School of Economics and Political Science

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and a number of other sources. It has five broad research areas:

1. Climate change impacts and resilience
2. Cutting emissions
3. Financing a better future
4. Global action
5. Protecting the environment

More information about the Grantham Research Institute is available at: www.lse.ac.uk/GranthamInstitute

Suggested citation:
Coppens L, Dietz S and Venmans F (2025) *Optimal climate policy under exogenous and endogenous technical change: making sense of the different approaches*. Grantham Research Institute on Climate Change and the Environment Working Paper 408. London: London School of Economics and Political Science

Optimal climate policy under exogenous and endogenous technical change: making sense of the different approaches *

Léo Coppens[†], Simon Dietz[‡], Frank Venmans[§]

March 5, 2025

Abstract

We analyse the large and diverse literature on technical change in integrated assessment models (IAMs) of climate change, with a view to understanding how different representations of technical change affect optimal climate policy. We first solve an analytical IAM that features several models of technical change from the literature, including exogenous technical change in abatement technologies, exogenous decarbonisation of the economy, endogenous technical change via learning-by-doing, and endogenous technical change via R&D (in particular, directed technical change). We show how these models of technical change impact optimal carbon prices, emissions and temperatures in often quite different ways. We then survey how technical change is currently represented in the main quantitative IAMs used to inform policy, demonstrating that a range of approaches are used. Exogenous technical change in abatement technologies and learning-by-doing are most popular, although the latter mechanism is only partially endogenous in some models. We go on to quantify technical change in these policy models using structural estimation, and simulate our analytical IAM numerically assessing the effect of technical change on optimal climate policy. We find large quantitative effects of technical change and large quantitative differences between different representations of technical change, both under cost-benefit and cost-effectiveness objectives.

Keywords: climate change; cost-benefit analysis; directed technical change; induced innovation; integrated assessment models; learning-by-doing; technical change

JEL codes: C61, O30, Q54, Q55, Q58

*Email for correspondence: f.venmans1@lse.ac.uk. This paper has benefited from valuable feedback from the editor, two anonymous referees, Aurelien Saussay, and participants at IAEE 2022, RES 2022, AERE 2022, EAERE 2022, IEW 2022, BEE 2022, the LSE Environmental Economics Workshop 2022, FAERE 2023, and a seminar at LSE. SD and FV gratefully acknowledge funding from the Grantham Foundation for the Protection of the Environment. FV acknowledges funding from EPSRC (Grant EP/W027887/1). LC acknowledges thesis funding from ICEDD ASBL. The authors have no conflicts of interest or financial interests to disclose.

[†]Warocqué School of Business and Economics, University of Mons, Place Warocqué 17, 7000 Mons, Belgium, leo.coppens@umons.ac.be

[‡]Department of Geography and Environment, Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, UK, s.dietz@lse.ac.uk

[§]Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, UK, f.venmans1@lse.ac.uk

1 Introduction

Integrated assessment models (IAMs) have played a central role in developing our understanding of economically efficient climate policies and continue to provide important inputs to decision-makers (e.g. IPCC, 2018, 2022). Technical change (hereafter TC) is one of the key assumptions in any IAM that estimates emissions abatement costs.

TC is a nebulous concept covering many mechanisms. Some are exogenous to abatement policy decisions. For example, there may be technological spillovers from non-climate R&D, such as general-purpose membrane technologies developed in the chemicals industry, which can also reduce the costs of clean hydrogen production. Some mechanisms are endogenous because they depend on climate R&D specifically. For example, in the current early stage of development of nuclear fusion, cost reductions depend primarily on R&D investment. Some mechanisms are endogenous to abatement policy because technology costs depend directly on deployment. For example, photovoltaic cells would still be expensive today if they had not been deployed at large scale. Usually, exogenous and endogenous TC will co-exist. For example, the development of Lithium-ion batteries for smartphones can reduce the future cost of electric-vehicle batteries (exogenous TC). However, the specific requirements of car batteries (e.g., large capacity and peak power) will be met more quickly if electric vehicles are produced at scale (endogenous TC).

IAMs need to make sense of the complex and diverse mechanisms of TC, simplify them and build appropriate model abstractions. Previous reviews have shown that modellers have taken many different approaches to TC (Löschel, 2002; Grubb et al., 2002; Sue Wing, 2006; Gillingham et al., 2008). At the same time, it is well known from other reviews that there is wide variation between IAMs in their estimated abatement costs of meeting pre-determined climate goals (Clarke et al., 2014; van Vuuren et al., 2020; Riahi et al., 2022), or their prescriptions of optimal warming (Gillingham et al., 2018). The missing piece of the puzzle is knowing what role TC plays in this variation. It is difficult to know because IAMs are rich and complicated, with many relevant differences. Therefore, it is not obvious how to construct a controlled comparison that leads to an understanding of the effect of different representations of TC on optimal climate policy. This is our aim in this paper.

We begin in Section 2 with a theoretical analysis of different TC models. Our fundamental aim is to sign the effect of TC on optimal climate policy, where we mostly use marginal abatement costs (MACs), carbon taxes, emissions and temperatures as sufficient statistics for climate policy (where relevant, we also discuss innovation subsidies of different kinds). We consider two kinds of climate policy, minimising the discounted sum of abatement costs and climate damages (cost-benefit analysis), and meeting a pre-determined climate constraint at minimum discounted abatement cost (cost-effectiveness analysis). Building on the foundation of a common set of equations for welfare, utility, warming and damages, we compare three different classes of TC model. The first is models of exogenous TC. We consider both exogenous TC in abatement technologies, which reduces abatement costs, and exogenous TC that reduces the emissions intensity of economic growth (so-called Autonomous Energy Efficiency Improvements or AEEI). The second class is models of endogenous TC based on learning-by-doing. In these models, the cost of abatement is a decreasing function of past cumulative abatement. The third class is models of endogenous TC based on R&D investment into abatement/clean technologies. We particularly focus on recent models of Directed Technical Change (DTC), which embed R&D in a model in which TC can improve either dirty or clean technologies.

The analysis delivers a rich set of results that we summarise in five Remarks, with derivations contained in the Appendices.

1. Exogenous TC that makes future abatement cheaper creates an incentive to abate less in the short run but more in the long run. Optimal carbon taxes are always lower, but emissions/temperatures are higher in the short run before ending up lower in the long run. Under cost-effectiveness, the long-run temperature is unaffected (it is constrained) but it is reached quicker.

2. Exogenous TC through AEEI has different effects. Its short-run effects are ambiguous, but plausibly it leads to lower optimal carbon taxes, emissions and temperatures. In the long run, however, it unambiguously leads to higher optimal carbon taxes, emissions and temperatures (again, under cost-effectiveness the long-run temperature is unaffected but it is reached quicker). This is because AEEI has two countervailing effects. On the one hand, it reduces business-as-usual (BAU) emissions, which makes lower emissions less costly to attain. On the other hand, by decarbonising the economy, it makes further emissions reductions more difficult. The balance of these effects is different at different stages of the low-carbon transition.

3. Endogenous TC based on learning-by-doing also has two countervailing effects. On the one hand, it makes future abatement cheaper, creating an incentive to wait, just like exogenous TC in abatement technologies. On the other hand, early abatement is what reduces future abatement costs, creating the opposite incentive to abate early. Optimal carbon taxes are lower, but an abatement subsidy is required to internalise the learning externality. The sum of these – the firm's MAC – can be higher or lower than without TC initially, so the effect on short-run emissions and temperatures is also ambiguous. In the long run, emissions and temperatures are lower, so emissions and temperatures under endogenous, learning-based TC can be lower all along the path. Under cost-effectiveness analysis, we show that the effect of endogenous, learning-based TC is qualitatively the same as exogenous TC in abatement technologies, but not as strong. The carbon budget is still used up more quickly and abatement is backloaded, but the endogenous future learning gain tempers the incentive to do so.

4. Models of endogenous TC based on R&D also create an incentive for early action, but not in the form of abatement. Instead, firms invest in R&D, incentivised by R&D subsidies, and delay abatement until R&D has reduced abatement costs. This means the optimal carbon tax, emissions and temperature trajectories look more like exogenous TC in abatement technologies than they do endogenous, learning-based TC. In fact, we show formally that they are the same as exogenous TC in abatement technologies, provided R&D investment costs are small relative to output.

5. We comment on clean technology dynamics in endogenous TC models. By design, learning-

based models are designed to replicate negative exponential experience curves that have been demonstrated empirically for a wide range of technologies (e.g., Way et al., 2022). To do this, they conform to the process of "fishing out", meaning that it becomes harder to find new ideas, the larger is the existing knowledge stock. By contrast, most DTC models conform to the opposite process of "standing on the shoulders of giants", meaning that existing knowledge makes new knowledge easier to develop. This logically means such models cannot reproduce experience curves.

Arguably a drawback of much of the literature is that endogenous, learning-based TC is represented at the exclusion of R&D, and *vice versa*. We finish Section 2 by proposing a model of endogenous R&D based on both.

Starting with Section 3, we shift our focus to quantitative models. This section contains a systematic survey of how TC is represented in the current crop of major quantitative IAMs, based on 22 families of models. We establish that a diversity of TC modelling approaches exists, as previous reviews also found. We find that TC is exogenous in the majority of models. Where it is endogenous, it is mostly based on learning-by-doing rather than R&D, a finding that is not obvious unless one-off models are excluded. Where models represent learning-by-doing, the incentive to abate early is often ignored in setting optimal policy, however. This leads us to identify a class of models where TC is 'semi-endogenous'.

Section 4 quantifies the effect of TC on climate policy. We develop a method of structural estimation, which enables us to calibrate the abatement cost and TC parameters of our model on the current crop of IAMs reviewed in the previous section. We use 'observed' variation in the timing of abatement and associated abatement costs across more than 700 IAM scenarios collected in two major databases (the Intergovernmental Panel on Climate Change or IPCC and the Network of Central Banks and Supervisors for Greening the Financial System or NGFS). This gives us an estimate of how much TC drives down abatement costs in current IAMs, supposing the process is either exogenous or endogenous. We then take this calibration of abatement costs/TC and apply it to a simple, quantitative IAM, whose structure builds on the theoretical analysis of Section 2. We numerically solve for optimal MACs, carbon prices and technology subsidies, as well as emissions and temperatures. We find that TC has a quantitatively large effect on the optimal paths under cost-benefit analysis, e.g., warming in 2100 is at least 0.5°C lower due to TC. We also find large differences between the two forms of TC considered, e.g., the difference in the initial MAC is \$73/tCO_2, attributable almost entirely to an optimal deployment subsidy to internalise the learning externality. Under cost-effectiveness analysis and an ambitious temperature constraint of 1.75°C, emissions and temperature paths are closer together due to the small available carbon budget, but the difference in MACs caused by TC is even larger.

Section 5 provides a discussion, focusing primarily on TC uncertainty.

Related literature

We connect to three main strands of literature. The first is papers investigating the effects of TC in IAMs of different varieties. This literature is large, diverse and diffuse across fields. One of our primary purposes in this paper is to survey it, so we provide citations throughout the following sections as our survey progresses.

Second, the model we build for theoretical and quantitative analysis draws on the recent literature on analytical IAMs (Golosov et al., 2014; Rezai and Van der Ploeg, 2016; van den Bijgaart et al., 2016; Dietz and Venmans, 2019; Traeger, 2023). As the moniker suggests, these models are intended to give analytical insights into the role of different parameters and assumptions, shining a light into the black box of richer quantitative IAMs. Our model has a similar level of complexity, uses specific representations from this literature, and delivers analytical as well as quantitative results.

Third, we contribute to the literature reviewing and synthesising quantitative IAMs, specifically those with a focus on emissions abatement (sometimes called energy models or energy system models). Many of the scenario runs of these models are summarised in IPCC reports (IPCC, 2022). Weyant (2017) and Nikas et al. (2019) provide recent, general overviews. Löschel (2002), Grubb et al. (2002), Sue Wing (2006) and Gillingham et al. (2008) are examples of earlier overviews of TC in these models, and van Vuuren et al. (2020) analyze the large differences in abatement costs between models. Relative to this literature, we make several contributions: we provide an up-to-date review; unlike previous reviews, we solve different TC representations analytically, which delivers a set of formal, precise comparisons; and we provide a method of isolating the quantitative effects of TC in these models on optimal paths, controlling for other factors.

2 Theoretical analysis

Throughout this section, we assume the economy is populated by a fixed number of households, which obtain utility from consumption of an aggregate good. The representative household's welfare is

$$W = \int_0^\infty e^{-\delta t} u\left(c(t)\right) dt, \tag{1}$$

where δ is the utility discount rate and $u\left(c(t)\right) = c(t)^{1-\eta}/(1-\eta)$, with η standing for the elasticity of marginal utility of consumption. In all the models we consider, households supply labour L inelastically and own capital K that is rented to firms.

Climate change is also modelled in the same way throughout. Global mean temperature T is linearly proportional to cumulative carbon dioxide emissions (Dietz and Venmans, 2019; Dietz et al., 2021):

$$T(t) = \zeta S(t). \tag{2}$$

Cumulative emissions $S(t) = \int P(t)dt$, where P stands for emissions. The slope of the relationship between temperature and cumulative emissions is governed by ζ, the so-called Transient Climate Response to cumulative carbon Emissions or TCRE (Collins et al., 2013).

Increasing temperatures cause climate damages, which only affect production of the aggregate good. Damages Ω are given by

$$\Omega(t) = \exp\left(-\frac{\gamma}{2} T(t)^2\right), \tag{3}$$

where γ is the damage function coefficient.

2.1 Models with an aggregate abatement cost function

The first models we look at link emissions directly with production of the aggregate consumer good:

$$P(t) = \sigma(t)\,[1 - \mu(t)]\,Y(t), \tag{4}$$

where σ is the carbon intensity of production Y and μ is emissions abatement expressed as a control rate. The total cost of abatement Λ is given by

$$\Lambda(t) = \exp\left(-\frac{\varphi(t)}{2}\mu(t)^2\right), \tag{5}$$

where φ is the slope of the MAC function.

The aggregate good is produced in a competitive market using capital and labour supplied by households. Production net of abatement costs and climate damages is

$$Y(t) = K(t)^{\alpha}\,[A_L(t)L(t)]^{1-\alpha}\,\Lambda(t)\Omega(t), \tag{6}$$

where A_L is labour productivity and α is the capital share.[1] Since climate damages and abatement costs enter the production function multiplicatively, they are both proportional to production. The number of firms is sufficiently large that climate damages are fully externalised.

Normalising the price of the aggregate good to one, the representative firm chooses capital and labour inputs to maximise profits:

$$\max_{K(t),L(t)} \pi(t) = K(t)^{\alpha}\,[A_L(t)L(t)]^{1-\alpha}\,\Lambda(t)\Omega(t) - r(t)K(t) - w(t)L(t) - \tau(t)P(t), \tag{7}$$

where r is the rental price of capital, w is the wage, and τ is a carbon tax that the firm may face. Any carbon tax revenue raised is returned to households lump-sum, so the household budget constraint implies that in competitive equilibrium

$$\dot{K}(t) = r(t)K(t) + w(t)L(t) + \Upsilon(t) - c(t) - \delta_K K(t), \tag{8}$$

where Υ represents the lump-sum transfers and δ_K is capital depreciation.

Exogenous TC

In many models, especially quantitative models but also analytical models focusing on aspects of the problem other than TC *per se*, TC is exogenous. Exogenous TC can be introduced to our model in two ways that are representative of the literature:

1. The slope of the MAC function φ decreases over time due to reductions in the cost of abatement technologies.

2. The carbon intensity coefficient σ decreases over time (as in, e.g., Nordhaus' DICE model). This decouples emissions growth from economic growth – economy-wide TC is emissions-

[1] The assumption here of Cobb-Douglas production reflects much of the literature and aids comparability with models below that explicitly represent energy inputs, but the same results can be obtained from a more general production function, as in the Appendices.

saving. It is often referred to as AEEI, although 'autonomous carbon efficiency improvement' would be more accurate.

If TC is exogenous, socially optimal emissions abatement is implemented using one policy instrument, a Pigouvian tax equal to the marginal damage cost of carbon emissions, a.k.a. the social cost of carbon (SCC).[2] In response, firms reduce emissions until their MAC is equal to the tax. Appendix A derives this well-known result from a more general model. For the model here, the optimality condition is:

$$\underbrace{\int_{u=t}^{\infty} e^{-r(u-t)} Y(u)\gamma\zeta S(u) du}_{\text{SCC}} = \tau(t) = \underbrace{\frac{\varphi(t)\mu(t)}{\sigma(t)}}_{\text{MAC}}, \tag{9}$$

where $r = \delta + \eta\dot{c}(t)/c(t)$ is the Ramsey rule for the consumption discount rate. The SCC is proportional to output, the damage function coefficient, the TCRE and cumulative future emissions, and it is decreasing in the discount rate.

The MAC is proportional to the MAC slope parameter φ. Hence, as TC decreases φ, it also decreases the MAC. The effect of this on optimal climate policy can be summarised as follows:

Remark 1 (Effects of exogenous TC in abatement technologies) *Compared to a model without TC, exogenous TC that decreases the slope of the MAC generates lower optimal MACs/carbon taxes, higher short-run emissions/temperatures but lower long-run emissions. Cumulatively, emissions are lower in the long run, as is warming.*

Appendix A provides the formal analysis in support of this. The basic intuition is that TC makes future abatement cheaper than today, and if TC is exogenous then all the social planner must do is wait for this to happen. This makes it optimal to abate emissions less in the short run but more in the long run. The net effect is not immediately obvious but in fact it must be lower cumulative emissions and warming in the long run. This is because TC makes the MAC at peak warming lower. Optimality then requires marginal damages to be lower accordingly. Since marginal damages are increasing in cumulative emissions and temperature, peak temperature must be lower.

This describes what happens if the social planner's objective is to maximise discounted net benefits. The planner's objective may instead be to minimise the discounted abatement costs of achieving a pre-determined cumulative emissions/temperature constraint. In that case, two things change. First, with a fixed emissions budget, the optimal carbon tax is not equal to the SCC but rather results from applying the Hotelling rule – it must increase at the interest rate, with an initial value equal to the present value of the MAC at zero emissions when the constraint binds. Second, if the constraint binds then evidently long-run temperatures are unaffected by TC. Exogenous TC that decreases abatement costs generates lower carbon taxes and higher short-run emissions as before, which now results in faster convergence to the constrained peak temperature. The social planner does not care about damages under cost-effectiveness and prefers to use up more of the cumulative emissions budget early on, anticipating cheaper abatement later.

[2]Or equivalently a cap on the quantity of emissions with an implied price equal to the social cost of carbon.

The effects of AEEI on optimal climate policy are different to falling abatement costs. AEEI exerts two contrasting forces on the optimal path. First, it reduces BAU emissions, which makes lower emissions less costly to attain. Second, by decarbonising the economy, it makes further emissions reductions more difficult. In an energy-efficient economy, a given relative reduction in emissions leads to lower absolute abatement ($d\mu(t) = -\sigma(t)Y(t)dP(t)$). These forces are illustrated by inspecting the effect of a change in σ on the MAC. Using (4) to substitute the emissions control rate out of the expression for the MAC, and differentiating with respect to σ,

$$\frac{\partial}{\partial\sigma(t)}\left(\frac{\varphi(t)}{\sigma(t)} - \frac{\varphi(t)P(t)}{\sigma(t)^2Y(t)}\right) = \frac{\varphi(t)}{\sigma(t)^2}\left(\frac{2P(t)}{\sigma(t)Y(t)} - 1\right).$$

A decrease in the carbon intensity of the economy decreases the MAC provided $2P(t) > \sigma(t)Y(t)$. $\sigma(t)Y(t)$ is simply BAU emissions. Therefore, this condition means that AEEI decreases the MAC when emissions are more than half of BAU, which if it is met[3] will be early in the transition, and increases the MAC when emissions are less than half of BAU, i.e., late in the transition. Early in the transition, the main effect of AEEI is to reduce BAU emissions, making lower absolute emissions easier to achieve. Late in the transition, the main effect of AEEI is to make further emissions reductions harder.

Appendix B derives optimal emissions, which by extension give optimal temperatures, under AEEI. We summarise the effects of AEEI on optimal climate policy as follows:

Remark 2 (Effects of exogenous TC: AEEI) *Compared to a model without TC, AEEI generates higher optimal MACs/carbon taxes, cumulative emissions and temperatures in the long run. The short-run effects of AEEI are ambiguous, but, assuming initial emissions are more than half of BAU, AEEI generates lower MACs/carbon taxes, emissions and temperatures.*

Since AEEI increases the MAC when emissions are less than half of BAU, AEEI increases the MAC at peak warming (zero emissions). Optimality then requires the SCC to be higher too. Since the SCC is an increasing function of temperature, peak temperature must be higher. This is the opposite of the effect of exogenous TC that decreases the slope of the MAC. In the short run, the effects of AEEI are ambiguous. As the effect on the MAC illustrates, the important condition is whether emissions are more or less than half of BAU. We might consider the former case more realistic in the short run, as large jumps in emissions are difficult. When emissions are more than half of BAU, AEEI reduces the MAC/carbon tax as shown above, and it reduces emissions and therefore temperatures.

If the planner's objective is cost-effectiveness and the temperature constraint binds, then AEEI does not affect peak warming. In principle, the effect of AEEI on the emissions path that uses up the cumulative emissions budget is ambiguous, but again if initial emissions are more than half of BAU, then AEEI decreases emissions in the short run, compensated by higher emissions in the long run and peak warming is reached later.

This summarises the effect of exogenous, economy-wide TC that is emissions-saving. Emission-saving TC is consistent with the declining trend in global emissions/output since the early 20th

[3]Technically, the optimal value of initial emissions can be anywhere between BAU and zero unless further constraints are applied. In our quantitative analysis, we introduce emissions inertia, which makes a big jump in initial emissions very costly.

century but it is not guaranteed. Indeed, the period before that was characterised by increasing emissions/output, plus in major developing countries emissions/output have been increasing until recently.[4] Logically, the opposite of our results would apply if exogenous, economy-wide TC were emissions-*using*. Di Maria and Smulders (2017) is one of few papers to explore similar cases, albeit in their model TC is endogenous.

Endogenous TC based on learning-by-doing

In other models, TC is endogenous because the cost of clean technologies decreases the more they are deployed (e.g., Goulder and Mathai, 2000; Van der Zwaan et al., 2002; Bramoullé and Olson, 2005; Grubb et al., 2024, and various quantitative/policy models reviewed in Section 3). This dynamic is often referred to as learning-by-doing. In the model here, it can be represented by writing the MAC as a function of cumulative abatement:[5]

$$\mathrm{MAC}(t) := \frac{\varphi\mu(t)}{\sigma(t)}\left(\frac{S_a(t)}{S_a(0)}\right)^{-\chi}, \tag{10}$$

where $S_a(t) = \int_{u=0}^{t} \sigma(u)\mu(u)Y(u)\,du$ is cumulative abatement and χ is the learning elasticity. Thus, for every percent increase in cumulative abatement, the MAC decreases by χ percent.

With a large number of firms, we assume for the sake of simplicity that the future learning gain from present abatement is fully externalised – incremental knowledge fully "spills over". Then, with a negative climate externality and a positive learning externality, the social planner generally needs two policy instruments, the carbon tax and an abatement subsidy s (e.g., a feed-in-tariff).[6] The firm's profit function becomes

$$\pi(t) = Y(t) - r(t)K(t) - w(t)L(t) - \tau(t)P(t) + s(t)\sigma(t)\mu(t)Y(t). \tag{11}$$

Lump-sum transfers in the household's budget constraint are now tax revenues net of subsidy payments.[7]

Socially optimal emissions abatement in this model is implemented by the following pricing rule:

$$\underbrace{\int_{u=t}^{\infty} e^{-r(u-t)}Y(u)\gamma\zeta S(u)du}_{\text{SCC}} + \underbrace{\int_{u=t}^{\infty} e^{-r(u-t)}Y(u)\,\frac{\varphi\chi}{2S_a(0)}\left(\frac{S_a(u)}{S_a(0)}\right)^{-\chi-1}\left[\sigma(u)\mu(u)Y(u)\right]^2 du}_{\text{Marginal future learning gain}} = \tau(t) + s(t). \tag{12}$$

Therefore, the social optimum can be achieved by setting the carbon tax equal to the SCC and the abatement subsidy equal to the marginal external gain in abatement knowledge, i.e., the

[4]Global Carbon Budget (2024), Bolt and van Zanden – Maddison Project Database 2023 – with major processing by Our World in Data (`https://ourworldindata.org/grapher/co2-intensity`).

[5]Appendix A develops a more general model of deployment-based endogenous TC.

[6]In the case where companies can appropriate part of the knowledge generated, the optimal subsidy will only cover the non-appropriable part.

[7]If lump-sum transfers are negative then taxation is assumed to be non-distortionary.

discounted marginal reduction in future abatement costs. However, a particular feature of this model of learning-based TC, where TC depends on cumulative aggregate abatement, is that the tax and subsidy are additive and the social optimum can in principle also be implemented by a single policy instrument, call it a 'learning-adjusted' tax, which is equal to the MAC. The intuition is that both climate damages and TC depend on cumulative abatement. However, we encourage readers to think of this tax/subsidy additivity result as more of theoretical curiosity. It does not hold in the model with clean and dirty energy discussed in the next section, or in other circumstances, such as when different abatement technologies have different learning rates (see Appendix A.3).

Remark 3 (Effect of learning-based TC) *Compared to a model without TC, endogenous TC based on learning-by-doing generates lower optimal carbon taxes when accompanied by abatement subsidies equal to the marginal future learning gain, but it has an ambiguous effect on emissions and temperatures in the short run. Emissions and temperatures are lower in the long run. Thus, emissions and temperatures can be lower all along the path.*

Appendix B provides the formal analysis in support of this. The basic intuition is that under endogenous TC from learning-by-doing, two countervailing effects are at play. First, there is the cost-reduction effect: TC makes future abatement cheaper than today, as with exogenous TC. Second, however, there is the endogenous future gain effect: the reason abatement is cheaper in the future is abatement today. This creates an incentive for early abatement and is optimally incentivised by an abatement subsidy (ignoring the additivity result above). Thus, while the optimal carbon tax is lower, the subsidy/marginal future learning gain can be large enough to produce higher abatement and lower emissions in the short run than without TC. In the long run, emissions and temperatures are lower, as they were in the previous case.

Under a cost-effectiveness objective, endogenous TC based on learning-by-doing has the same effect as exogenous TC in abatement technologies – lower carbon taxes, higher short-run emissions and consequently faster convergence to the constrained peak temperature – but the positive effect on emissions and temperatures in the short run is not as strong as under exogenous TC. Again, the difference is due to the endogenous future learning gain, which tempers the incentive to backload abatement.

Price-induced TC

In a few models, technology is directly linked to prices (e.g., Jakeman et al., 2004; Wilkerson et al., 2013). We can connect the analysis up to this point with such models.

The key driver of endogenous TC in the above model is a state variable (cumulative emissions S_a), which can be interpreted as a knowledge stock. The knowledge stock's equation of motion depends on the quantity of abatement. In general, we can write $\dot{S}_a(t) = f(\mu(t))$. Suppose the knowledge stock is a function of the carbon price rather than abatement quantity. This can give the same dynamics, because any quantity of abatement has a corresponding carbon price: $\mu(t) = g(\tau(t))$, where $g(\cdot)$ is the inverse MAC function. As a result, the same equation of motion can now be written as a function of τ rather than μ: $\dot{S}_a(t) = f(g(\tau(t)))$. Integrating this equation gives current technology/knowledge, which will be a function of *past* carbon prices.

Remark 3 thus applies to this kind of model. For models in which abatement affects other prices, which in turn affect the build up of knowledge (via the function $f(\cdot)$), abatement indirectly affects the dynamics of the knowledge stock, and again we can say in general that there will be an extra marginal learning gain term that drives a wedge between the SCC and the MAC as in (12).

Technology parameters may be a function of the current price rather than past prices (Jorgenson and Wilcoxen, 1993), and technology is not modelled as a stock variable but rather as an instantaneous variable. As a result, TC is not path-dependent (there is no knowledge stock) and it makes more sense to consider it as a model of exogenous TC (Gillingham et al., 2008).

2.2 Models with clean and dirty energy substitution

Endogenous TC based on learning-by-doing

Instead of modelling emissions as a function of aggregate output and emissions abatement via an aggregate abatement cost function, many models explicitly link emissions with the use of dirty/fossil energy (e.g., Golosov et al., 2014; Fried, 2018; Hart, 2019; Hassler et al., 2021, and various quantitative/policy models reviewed in Section 3). Taking the model of endogenous TC based on learning-by-doing, here we introduce clean and dirty energy intermediates in production. Final-good production is now

$$Y(t) = K(t)^{\alpha} E(t)^{\nu} \left[A_L(t)L(t)\right]^{1-\alpha-\nu} \Omega(t), \tag{13}$$

where E is an energy composite and ν is the energy expenditure share. As the final-good production function is Cobb-Douglas, this is close to the model in Golosov et al. (2014). Other models use a constant elasticity of substitution (CES) function for final-good production (e.g., Hassler et al., 2021).

The composite, intermediate energy good is produced by combining clean and dirty energy, assuming a constant elasticity of substitution (CES):

$$E(t) = \left[E_c(t)^{\epsilon} + E_d(t)^{\epsilon}\right]^{1/\epsilon} \tag{14}$$

where ϵ determines the elasticity of substitution between clean and dirty energy (indexed c and d respectively). How clean and dirty energy are produced varies between models, with some using labour (e.g., Fried, 2018; Golosov et al., 2014) and others using final goods as the input (Hart, 2019). This choice is not consequential for our purposes. We follow the latter approach, so $E_i(t) = A_i(t)X_i(t)$, $i \in \{c, d\}$, where A_i is clean/dirty energy productivity and X_i is the quantity of final goods used in clean/dirty energy production.[8] Fossil fuels are assumed to be abundant (i.e., specifically coal is abundant). Expressing dirty energy use in tonnes of carbon dioxide, emissions are identified with dirty energy use, $P = E_d$. A caveat is that with CES between clean and dirty energy, the MAC is implicitly infinite at zero emissions. Therefore, in

[8] Technically X_d is the quantity of final goods used in dirty energy production *adjusted for* the productivity of fossil-fuel extraction, as explained by Hart (2019). Dirty energy production $E_d(t) = A_d(t)D_d(t)$, where D_d stands for fossil inputs, and $D_d(t) = A_{dx}X_d(t)$ is fossil-fuel extraction with extraction productivity A_{dx}. The price of the final good is one and productivity is also constant since the final good is unchanging, so unit extraction costs are constant and equal to $1/A_{dx}$.

numerical applications this structure would need to be augmented, for example, with a backstop technology if the SCC is large enough to make zero or negative emissions desirable.

In general, learning-by-doing applies to both clean and dirty energy. The experience-curve dynamic can be produced by writing clean/dirty energy productivity as a function of a stock of knowledge:

$$A_i(t) = A_i(0)\left(\frac{H_i(t)}{H_i(0)}\right)^{\chi}, \tag{15}$$

where H is cumulative aggregate knowledge, which accumulates in proportion to deployment/use of the respective energy sources:

$$\dot{H}_i(t) = E_i(t). \tag{16}$$

This formulation abstracts from knowledge depreciation, which should be slow. Knowledge is assumed to be technology-specific and there are no spillovers from one technology to the other. On the contrary, using green technology will slow down the knowledge built up in the dirty sector and *vice versa*, which is a form of crowding out.

We remain in a competitive market and we assume the planner can use two policy instruments, a carbon tax and a clean-energy subsidy. The profit function of the firm in this case is

$$\pi(t) = Y(t) - r(t)K(t) - w(t)L(t) - [p_c(t) - s(t)]\,E_c(t) - [p_d(t) + \tau(t)]\,E_d(t). \tag{17}$$

Socially optimal emissions abatement in this model can be implemented by setting the carbon tax equal to the SCC *minus* the marginal future learning gain from dirty energy use, and the clean-energy subsidy equal to the marginal future learning gain from clean energy use:

$$\underbrace{\int_{u=t}^{\infty} e^{-r(u-t)}Y(u)\gamma\zeta S(u)du}_{\text{SCC}} - \underbrace{\int_{u=t}^{\infty} e^{-r(u-t)}Y(u)\,[E_d(u)/E(u)]^{\epsilon}\,\frac{\chi}{H_d(u)}du}_{\text{Marginal learning gain: dirty}} = \tau(t), \tag{18}$$

$$\underbrace{\int_{u=t}^{\infty} e^{-r(u-t)}Y(u)\,[E_c(u)/E(u)]^{\epsilon}\,\frac{\chi}{H_c(u)}du}_{\text{Marginal learning gain: clean}} = s(t). \tag{19}$$

The firm will use dirty energy up to the point where its marginal product equals the price of dirty energy plus the tax,

$$\frac{\partial Y(t)}{\partial E_d(t)} = [p_c(t) - s(t)]\left(\frac{A_d(t)}{A_c(t)}\right)^{\epsilon}\left(\frac{E_d(t)}{E_c(t)}\right)^{\epsilon-1} = p_d(t) + \tau(t). \tag{20}$$

Similarly, it will use clean energy up to the point where its marginal product equals the price of clean energy minus the subsidy.

The effect of TC in this model can be seen by decomposing the productivity term in (20):

$$\left(\frac{A_d(t)}{A_c(t)}\right)^{\epsilon} = \left(\frac{A_d(0)}{A_c(0)}\right)^{\epsilon}\left(\frac{H_c(t)}{H_c(0)}\right)^{\epsilon\chi}\left(\frac{H_d(t)}{H_d(0)}\right)^{-\epsilon\chi}. \tag{21}$$

This decreases in cumulative clean energy use with an elasticity of $-\epsilon\chi$ and increases in cumulative dirty energy use with an elasticity of $\epsilon\chi$.

As mentioned above, the carbon tax and the innovation subsidy are not perfect substitutes in this model. A carbon tax will lead to both less energy use and energy substitution from dirty to clean. Although a clean-energy subsidy will also incentivise substitution from dirty to clean, by itself it will lead to more energy use, not less. Yet, defining the MAC as the difference between the marginal products of dirty and clean energy, MAC $:= dY(t)/dE_d(t) - p_d(t) - dY(t)/dE_c(t) + p_c(t) = \tau(t) + s(t)$. So, although the tax and the subsidy are not perfect substitutes, the MAC is equal to their sum.

With $0 < \chi < 1$, there are decreasing marginal returns to the energy knowledge stocks, reflecting the empirical observation that learning-by-doing is slower in mature technologies. Therefore, with so much historical use of dirty energy, marginal learning gains from further dirty energy use will be small, and so will the adjustment to the optimal carbon tax. In fact, the vigilant reader will notice that in principle the planner could use three policy instruments: the carbon tax and the clean-energy subsidy, plus a dirty-energy subsidy equal to the marginal future learning gain from dirty energy use. However, the dirty-energy subsidy is additive to the carbon tax and it makes little sense to implement it separately, especially given its likely small size.

These additional features aside, Remark 3 and its extension to cost-effectiveness also apply to the effects of TC on optimal climate policy in this case.

Endogenous TC based on R&D: directed technical change

So far, TC has either been exogenous or it has been endogenous because of learning-by-doing. The third important class of model makes TC endogenous because future abatement costs depend on current R&D investments. Early models of this type included Goulder and Mathai (2000); Smulders and De Nooij (2003); Buonanno et al. (2003); Popp (2004), with diverse mechanisms through which R&D reduces abatement costs. Within the last decade, R&D-based TC has been popularised by applying models of directed technical change (DTC: e.g., Acemoglu et al., 2012, 2016; Bretschger et al., 2011; Fried, 2018; Hart, 2019; Hassler et al., 2021). Distinguishing between clean and dirty energy is the most obvious (and popular) way to introduce a direction to TC.

We start with a general functional form for R&D-based TC:

$$\dot{A}_i(t) = \psi\left(z_i(t), A_i(t)\right), \tag{22}$$

where z_i stands for R&D investments. For $\partial\psi/\partial A_i > 0$, we have so-called "standing on the shoulders of giants", meaning that existing knowledge makes new knowledge easier to develop. The opposite case, $\partial\psi/\partial A_i < 0$, is called "fishing out", meaning that it becomes harder to find new ideas, the larger is the existing knowledge stock. We will return to this important issue below.

A feature of R&D models is usually imperfect competition – the prospect of a monopoly achieved by patenting a new invention incentivises costly R&D. This complicates the model dynamics, so DTC models, which are written in discrete time, typically assume the patent lasts just one period (e.g., Acemoglu et al., 2012). We follow this assumption in our continuous-time model, noting the assumption is less realistic when the time step is infinitesimal. Practically, the implication is that R&D needs to be subsidised in order for it to happen, because the monopoly vanishes instantaneously as an approximation for vanishing in a few years as it does in discrete-time models. If we made the model more complex by introducing an explicit time dependence of monopoly profits, the main addition would be that some R&D happens without subsidies. Appendix D discusses such a model with a private knowledge stock that gradually becomes public. The rest of the model is the same as the previous section. We assume that R&D costs are homogeneous of degree one.

The social planner has three tools: a carbon tax, and subsidies for clean- and dirty-energy R&D. The carbon tax and the dirty-energy R&D subsidy are not additive in this case. That is because TC no longer depends on cumulative abatement. The firm's profit function becomes

$$\begin{aligned}\pi(t) = Y(t) - r(t)K(t) - w(t)L(t) - p_c E_c - \left[p_d(t) + \tau(t)\right] E_d(t) \\ -c(z_c(t) + z_d(t)) + s_c(t)z_c(t) + s_d(t)z_d(t),\end{aligned} \tag{23}$$

where $c(\cdot)$ is an R&D cost function. This is where crowding out happens. Most models assume costs are linear in $z_c + z_d$, where z is often described as the number of scientists in a given sector.[9] If the number of scientists is assumed fixed (e.g., Acemoglu et al., 2012), crowding out of dirty innovation by clean is strong because additional scientists cannot be recruited from outside the R&D sector. In Popp (2004), which is an R&D model but not a DTC model *per se*,[10] crowding out is also strong because the opportunity cost of innovation is set at three times the investment in R&D.

Socially optimal emissions abatement in this model is implemented by setting the carbon tax equal to the SCC and the clean- and dirty-energy subsidies equal to the marginal productivity benefits of R&D in the respective sectors:

$$\underbrace{\int_{u=t}^{\infty} e^{-r(u-t)} Y(u)\gamma\zeta S(u)du}_{\text{SCC}} = \tau(t), \tag{24}$$

$$\underbrace{\int_{v=t}^{\infty} \left[\exp\left[-\int_{u=t}^{v} \left[r(u) + \partial\psi/\partial A_c(u)\right] du\right] \partial Y(v)/\partial A_c(v) dv\right] \partial\psi/\partial z_c(t)}_{\text{Marginal benefit of clean R\&D}} = s_c(t), \tag{25}$$

$$\underbrace{\int_{v=t}^{\infty} \left[\exp\left[-\int_{u=t}^{v} \left[r(u) + \partial\psi/\partial A_d(u)\right] du\right] \partial Y(v)/\partial A_d(v) dv\right] \partial\psi/\partial z_d(t)}_{\text{Marginal benefit of dirty R\&D}} = s_d(t). \tag{26}$$

Firms will respond by reducing emissions until their MAC equals the tax and by investing in R&D until marginal investment costs in each sector equal the subsidies.

In this model, TC is endogenous in the sense that it is affected by the future trajectory of the carbon price. This can be seen from (25) and (26), because the marginal benefit of energy R&D is a function of the marginal product of energy knowledge, and this in turn depends on energy use, which is affected by the carbon price. However, this does not change the fact that in R&D/DTC models, the MAC equals the SCC (or the MAC follows the Hotelling rule under

[9] Often the innovation sector is represented as a separate sector, which receives a patent when a successful innovation is realised. We have assumed innovation is realised in-house, but the dynamics are the same.

[10] Although Popp has a clean and dirty sector, TC in the dirty sector is exogenous.

cost-effectiveness analysis). Furthermore, at the optimum a model of exogenous TC via falling abatement costs will have the same dynamics as this model of R&D-based TC:

Remark 4 (Effect of R&D-based TC) *Replace a model of endogenous TC based on R&D with a model of exogenous TC, ensuring abatement costs follow an identical time path. The models will have the same optimal carbon taxes, emissions and temperatures provided R&D investment costs are small relative to output.*

Appendix E.3 shows this result formally. The basic intuition is that early action is optimal in R&D models, but unlike learning-based models the early action comes in the form of clean energy R&D investment, rather than abatement. The planner optimally invests in R&D before deploying clean technology, with effects on optimal carbon taxes, emissions and temperatures that with suitable calibration are identical to when the planner waits for technology costs to fall exogenously. This result depends on the assumption that R&D investment costs are small relative to output/consumption growth, because these costs reduce consumption and therefore have the potential to alter marginal climate damages and the discount rate, i.e., to alter the SCC. It is highly likely that R&D investment costs are relatively small (IPCC, 2022). It is also important to stress that although Remark 4 holds at the optimum, the exogenous TC model so calibrated will not give the same climate-policy dynamics away from the optimum. Therefore, the result does not imply exogenous TC will always capture the dynamics of endogenous, R&D-based TC. It is merely intended to sharpen our intuitions about how optimal policies differ qualitatively.

Technology cost dynamics in learning- versus R&D-based models

We now return to standing on the shoulders of giants versus fishing out, and the dynamics of TC in learning- versus R&D-based models. Besides having different implications for optimal abatement, these two classes of endogenous TC model tend to predict different technology-cost trajectories, which in turn can lead to very different optimal carbon taxes. For example, Acemoglu et al. (2012) famously found that optimal carbon taxes could be very low because a temporary clean R&D subsidy was so effective in reducing future abatement costs.

First, note that we can rewrite the learning-based model using energy productivity A_i instead of knowledge H_i as the state variable. Taking the time derivative of (15) gives us the new equation of motion:

$$\dot{A}_i(t) = \frac{A_i(0)^{\frac{1}{\chi}}}{H_i(0)} A_i(t)^{1-\frac{1}{\chi}} E_i(t). \tag{27}$$

In this equation, instantaneous productivity on the right-hand side is raised to the power $1-1/\chi$. Assuming $0 < \chi < 1$, this power is negative, consistent with $\partial\psi/\partial A_i < 0$ and giving us fishing out.

This is in contrast to most DTC models. Take Acemoglu et al. (2012), in which the equation of motion of clean/dirty energy productivity, written in continuous time, is

$$\dot{A}_i(t) = \xi_i z_i(t) A_i(t), \tag{28}$$

where ξ_i is the expected growth rate of productivity. Because productivity growth is proportional to the current level of productivity, it is clear that in this model there is a strong standing on the shoulders of giants effect.

Later DTC models such as Fried (2018) introduce additional features that weaken the size of the standing on the shoulders of giants effect, but it still holds. In Fried (2018),

$$\dot{A}_i(t) = \xi_i z_i(t)^\iota A_i(t) \left[\frac{A_i(t)}{A_i(t) + A_j(t)}\right]^{-\varsigma} = \xi_i z_i(t)^\iota A_i(t)^{1-\varsigma} \left[A_i(t) + A_j(t)\right]^\varsigma , \qquad (29)$$

where $\iota \in (0,1)$ implies diminishing returns to R&D within a period (the "stepping on toes effect"), creating an incentive to smooth out R&D effort over time, and the factor $\left[\frac{A_i(t)}{A_i(t)+A_j(t)}\right]^{-\varsigma}$ introduces cross-sectoral technology spillovers between sectors. Fried assumes $\varsigma \in (0,1)$. Thus, cross-sectoral spillovers reduce the elasticity of productivity growth with respect to current productivity, but still $\partial\psi/\partial A_i > 0$ – there is standing on the shoulders of giants. It is perhaps worth emphasising that despite the superficial similarity of language, standing on the shoulders of giants and stepping on toes are mutually compatible dynamics. The former amounts to $\partial\psi/\partial A_i < 0$, the latter to $\partial^2\psi/\partial {z_i}^2 < 0$.

An exception is the DTC model of Hart (2019). In Hart's model, the equation of motion of the knowledge stock is

$$\dot{H}_i(t) = -\delta_H H_i(t) + A_L(t)^{\varsigma_1} \left(H_i(t) + \varsigma_2 H_j(t)\right)^{1-\varsigma_1} \xi z_i(t)^{1-\iota}. \qquad (30)$$

The first term is knowledge depreciation and the second term is knowledge accumulation. Within the second term, $H_i + \varsigma_2 H_j$ ensures there is standing on the shoulders of giants overall in this equation, albeit with diminishing returns given Hart's assumption that $\varsigma_1 \in (0,1)$ and further attenuation from positive cross-sectoral spillovers ($\varsigma_2 \in (0,1)$). Again, $\iota \in (0,1]$ produces the stepping on toes effect in relation to instantaneous R&D effort.

However, Hart's model separates productivity from knowledge, which is unusual in endogenous TC models. The justification Hart gives is that energy productivity has a physical maximum. This function is concave, so the marginal gains from accumulating knowledge are decreasing, which is important. Concretely, energy productivity evolves according to

$$A_i(t) = \bar{A}_i \frac{H_i(t)}{H^* + H_i(t)}, \qquad (31)$$

where $\bar{A}_i$ is the theoretical maximum productivity as knowledge H_i goes to infinity and H^* is the knowledge stock at which productivity is $\bar{A}_i/2$.

Combining Equations (30) and (31), we get

$$\dot{A}_i(t) = \frac{\left[\bar{A}_i - A_i(t)\right]^2}{\bar{A}_i H^*} \left[-\delta_H \frac{A_i(t)H^*}{\bar{A}_i - A_i(t)} + A_L(t)^{\varsigma_1} \left[\frac{A_i(t)H^*}{\bar{A}_i - A_i(t)} + \varsigma_2 H_j(t)\right]^{1-\varsigma_1} \xi z_i(t)^{1-\iota}\right]. \qquad (32)$$

It can be shown that the sign of the derivative of this with respect to A_i is positive for small A_i, but for large A_i the knowledge decay term comes to dominate and the derivative is negative. Thus, Hart's model has standing on the shoulders of giants at low productivity levels, but fishing out at high productivity levels.

Remark 5 (Comparative dynamics of endogenous TC) *Learning-based models of TC exhibit fishing out whereas most DTC models exhibit standing on the shoulders of giants. Breaking the link between knowledge and productivity allows DTC models to exhibit fishing out when technologies are mature.*

The importance of this remark lies in the fact that learning-based models are designed to fit the empirical regularity of Wright's law and negative exponential experience curves. To reproduce such an experience curve, there must be diminishing marginal returns to R&D as the knowledge stock increases, but early DTC models have increasing returns. The standing on the shoulders of giants effect is motivated by trends in patent counts. But both effects can co-exist if the marginal productivity of patents is decreasing. One way to capture both of these effects is to introduce diminishing marginal energy productivity gains with respect to knowledge, which is what happens in Hart's model at high levels of knowledge.

Endogenous TC based on learning *and* R&D

The idea that knowledge only accumulates through learning-by-doing or R&D is unrealistic, compared with the idea that it responds to both (Grubb et al., 2021). Unfortunately, it is a feature of the literature that models tend to use one representation at the expense of the other. But they can be combined. Represent the interdependency between deployment and R&D by an equation of motion of the general form

$$\dot{A}_i(t) = \psi\left(z_i(t), E_i(t), A_i(t)\right). \tag{33}$$

The optimal policy now requires a deployment subsidy for clean energy

$$\underbrace{\int_{v=t}^{\infty}\left[\exp\left[-\int_{u=t}^{v}\left[r(u)+\partial\psi/\partial A_c(u)\right]du\right]\partial Y(v)/\partial A_c(v)dv\right]\partial\psi/\partial E_c(t)}_{\text{Marginal innovation benefit of clean deployment}} = s_c^{deploy}(t), \tag{34}$$

on top of the R&D subsidy for clean energy from Eq. (25). Technically, the same goes for dirty energy.

Note that the integral in the formula corresponds to the forward-looking value of an extra unit of $A_c(t)$ and is also an element of the R&D subsidy. This leads to the useful result that the ratio of the deployment subsidy and the R&D subsidy should be equal to the ratio of their marginal productivities:

$$\frac{s_c^{deploy}}{s_c^{R\&D}} = \frac{\partial\psi/\partial E_c(t)}{\partial\psi/\partial z_c(t)}. \tag{35}$$

The right-hand side may be amenable to empirical estimation. If so, this may provide a guide to the elusive question of how much to subsidise R&D versus deployment **?**.

The MAC defined as the difference in net productivity between dirty and clean energy is higher than the SCC due to the effect of deployment on innovation:

$$MAC = SCC - s_d^{deploy} + s_c^{deploy}. \tag{36}$$

Thus, this model will have optimal paths qualitatively similar to endogenous, learning-based TC. Typically, the clean-energy deployment subsidy will dominate the dirty-energy deployment subsidy due to the larger innovation potential of clean energy. Note that this increase in the MAC is still present in a model where there is no innovation without R&D effort, for example, in a model with the following equation of motion: $\dot{A}_i(t) = z_i(t) + \rho z_i(t) E_i(t)$.

3 Survey of quantitative IAMs

We now move from theory models to quantitative IAMs intended to inform policy. This section describes our systematic review of how TC is represented in these models currently.

We first compiled a list of candidate IAMs, populating the list using a set of international databases/web resources and previous reviews on the topic.[11] The resulting long list comprised 87 models. We then screened this long list of IAMs based on the following criteria for inclusion. First, the model must be global. Second, the model must be in current/recent use, which we defined as having yielded a publication within the three years prior to undertaking our review. Third, the model must have been designed to estimate mitigation costs from the energy system (this excluded models primarily intended to estimate the SCC, and it also excluded specialist land-use models). Fourth, the model must have been used in multiple papers or projects (we excluded 'one-off' models). Lastly, we consolidated versions of the same model into a single family. After screening and combining, we were left with 22 model families for analysis of their representation of TC. These are listed in Table 1.

The table describes the type of model and then classifies the models' representation of TC. We divide the models into exogenous, endogenous and what we call 'semi-endogenous' TC, and for endogenous TC models we further specify whether the process is based on learning or R&D. We use semi-endogenous to describe models in which current deployment of an abatement technology makes the technology cheaper in the future, but where this learning mechanism does not affect the optimal price/subsidy trajectory. In other words, models with semi-endogenous TC include larger future deployment as technologies become cheaper (learning-by-doing), but omit the incentive for early abatement anticipating the endogenous future gain effect.[12]

The results of our systematic review are as follows. First, the diversity of modelling approaches to TC identified in earlier reviews endures today. Second, we find that TC is exogenous in the majority of models. Four models include endogenous TC, and five other models have semi-endogenous TC. TC is exogenous in the remaining 13 models. A pre-requisite for fully endogenous TC is the ability to optimise the model intertemporally; only eight models do

[11] In particular, we used the IPCC AR6 Scenario Explorer and Database hosted by IIASA (`https://data.ece.iiasa.ac.at/ar6/#/workspaces`); the web resources of the Integrated Assessment Modeling Consortium or IAMC (`https://www.iamconsortium.org/resources/models-documentation/`), the United Nations Framework Convention on Climate Change or UNFCCC response measures modelling tools (`https://unfccc.int/topics/mitigation/workstreams/response-measures/modelling-tools-to-assess-the-impact-of-the-implementation-of-response-measures`), the Stanford University Energy Modeling Forum (Böhringer et al., 2021), and previous review articles by Gillingham et al. (2008) and Nikas et al. (2019).

[12] To ensure our classification of models as semi-endogenous was reasonable, we contacted the relevant modelling teams to explain our concept of semi-endogenous TC and check our characterisation of their model. We contacted eight modelling teams (IMAGE, GTEM, POLES, E3ME, GEM-E3, EPPA, IMACLIM-R, IGEM) and received answers from seven of them.

this, as column 3 shows. Lastly, of the models with semi-endogenous or endogenous TC, the majority represent learning-by-doing, with R&D investments explicitly represented in only two models. In their sample, Gillingham et al. (2008) identified a larger share of models with R&D, but most of these models were one-off developments: R&D is less prevalent as a TC mechanism in the most commonly used core versions of IAMs that feed into inter-comparison exercises like the IPCC scenario database.

4 Quantitative analysis

The purpose of this section is to estimate the effects of TC in the quantitative/policy IAMs on optimal climate policy. This is not easy. Previous literature asking similar questions has relied on regression of IAM outputs on model features (e.g. Kuik et al., 2009), sensitivity analysis using a single model (e.g. Manne and Richels, 2004), and harmonised runs of multiple IAMs (e.g. Gillingham et al., 2018). Each approach has its pros and cons. Standard, regression-based meta-analysis provides a convenient way to explore a wide range of model features, but it is typically limited to qualitative results on the effects of TC (since model features are represented as dummy variables). Given the nature of the data, regression also faces identification problems such as low statistical power and multi-collinearity. Sensitivity analysis with individual IAMs permits a tightly controlled experiment into the effects of TC but is limited to an individual model structure. By contrast, harmonised runs of multiple IAMs allow model uncertainty to be explored, but it is practically difficult to evaluate many models this way, let alone multiple parameterisations of each model.

Our approach consists of two steps. The first step is structural estimation of how much TC reduces abatement costs in the current crop of IAMs. The second step is to build a simple IAM, based on the theory models of Section 2, and use it to simulate optimal climate policy based on the TC parameter estimates from the first step. Thus, our approach is most similar to running sensitivity analysis with an individual IAM, but by ensuring through the first step that the model replicates TC in more complex IAMs, and by using a simple model structure broadly in the tradition of analytical IAMs, the model is set up to perform more of a meta-analytical function.

We implement both exogenous TC in abatement technologies and endogenous, learning-based TC, as these are the two principal types of TC that feature in the quantitative/policy IAMs. R&D-based TC is not included in this analysis, because it rarely features in quantitative/policy IAMs and when it does so it is alongside learning-based TC – we cannot differentiate their effects in the data. However, we know from Section 2 that R&D-based TC will give approximately the same optimal climate policy paths as exogenous TC when the metrics of interest are carbon taxes, emissions and temperatures, as long as R&D costs are not large enough to shift the consumption growth rate. We run our simple IAM under cost-benefit and cost-effectiveness objectives.

Model	Type	Intertemporal optimisation	Exgenous or endogenous TC	LbD/R&D	References
AIM/CGE	CGE	No	exogenous		Fujimori et al. (2017a,b)
DICE	Optimal growth	Yes	exogenous		Nordhaus (2017)
DNE21+	Energy System	Yes	exogenous*		Sano et al. (2006); Wada et al. (2012)
E3ME	Macroeconometric	No	semi-endogenous	LbD/R&D	Mercure et al. (2018) Cambridge Econometrics (2019)
ENV-Linkages	CGE	No	exogenous		Château et al. (2014)
EPPA	CGE	No	exogenous		Chen et al. (2015); Jacoby et al. (2006) Octaviano et al. (2016)
GCAM	Other IAMs	No	exogenous		Bond-Lamberty et al. (2022); Calvin et al. (2017)
GEM-E3	CGE	No	semi-endogenous	LbD	Capros et al. (2013)
GTAP-E	CGE	No	exogenous		Burniaux and Truong (2002); Corong et al. (2017)
GTEM	CGE	No	semi-endogenous	LbD	Jakeman et al. (2004); Pant (2007)
ICES	CGE	No	exogenous*		Eboli et al. (2010); Parrado and De Cian (2014)
IGEM	CGE	Yes	exogenous		Goettle et al. (2007)
IMACLIM-R	CGE	No	semi-endogenous	LbD	Bibas et al. (2022)
IMAGE	Other IAMs	Yes	endogenous	LbD	Stehfest et al. (2014)
MARKAL/ TIMES	Energy System	Yes	endogenous**	LbD	Loulou et al. (2016)
MESSAGEix-GLOBIOM	Other IAMs	Yes	exogenous*		Fricko et al. (2017); Krey et al. (2020) Messner (1997)
PACE	CGE	Yes***	exogenous		Böhringer et al. (2009a); Gavard et al. (2022)
Phoenix	CGE	No	exogenous		Sue Wing et al. (2011); Lucena et al. (2018)
POLES	Energy System	No	semi-endogenous	LbD	Keramidas et al. (2017)
REMIND	Optimal growth	Yes	endogenous	LbD	Luderer et al. (2015)
WEM	Energy System	No	exogenous		IEA (2021)
WITCH	Optimal growth	Yes	endogenous	LbD/R&D	Emmerling et al. (2016)

* A study adding endogenous change exists, but this is not incorporated in the main model (e.g., Messner (1997) for MESSAGE; Parrado and De Cian (2014) for ICES; Sano et al. (2006) for DNE21).
** Most applications do not use the endogenous TC feature of the model.
*** There is an extension allowing intertemporal optimisation of the model (Böhringer et al. (2009b)).

Table 1: Representation of technical change in 22 IAMs. In models with learning-by-doing (LbD), the cost of a technology depends on past (cumulative) deployment. In models with R&D, the cost of a technology depends on past investments in R&D. The "Other IAMs" category contains Energy System models coupled with other modules such as land use.

4.1 Structural estimation of TC

Our structural estimation uses 'observed' variation in the timing of abatement and associated abatement costs across a large number of IAM scenarios collected in two major databases (the IPCC and the NGFS).[13] This gives us an estimate of how much TC drives down abatement costs in current IAMs, supposing the process is either exogenous or endogenous, learning-based.

The abatement cost/TC parameters are calibrated on 739 scenarios from 13 leading IAM families,[14] obtained by pooling results from the IPCC and NGFS databases. We exploit variation between IAM scenarios in total abatement costs, MACs and emissions. If the underlying IAMs were static, a given quantity of abatement would cost the same whenever it happens. But with TC, a given quantity of abatement is more costly the earlier it happens. This is our identifying variation. A model of exogenous TC can be estimated by assuming that observed cost reductions in the dataset are driven by time. Alternatively, a model of endogenous, learning-based TC can be estimated by assuming that observed cost reductions are a function of cumulative abatement. We cannot estimate a mixed exogenous/endogenous model – the underlying models are either exogenous or endogenous, and there is high collinearity between time and cumulative abatement. However, the parameter estimates recovered from the pure exogenous and pure endogenous, learning-based TC models could still be used in a mixed TC model, where the MAC function is a weighted average of the two.

Parameter estimation is made using the Generalised Method of Moments, estimating total and marginal abatement cost functions simultaneously. This allows us to obtain more robust results: although the MAC function is more economically meaningful as it determines the FOCs of the optimum, the MAC functions of the underlying IAMs could be non-linear. We give equal weight to the errors of both the total and marginal abatement cost functions and assume they are independent and normally distributed. For both exogenous and endogenous TC, we report four specifications. First, we fit the exogenous or endogenous TC parameters on all the IAMs in the databases. This utilises all the available variation between models, but treats IAMs with endogenous, learning-based TC *as if* they had exogenous TC, and *vice versa*. The second specification fits the exogenous TC parameters only on the subset of IAMs with exogenous TC, and similarly for endogenous, learning-based TC. Third, we include a fixed-effects specification, where we allow each model to have its own initial MAC slope.[15] Fourth, we include a fixed-effects specification where each model has its own initial MAC slope *and* its own TC rate. The fixed effects control for unobserved, model-specific factors, which could confound the effects of TC. The fourth specification is the most flexible and is the one we use. Table 2 summarises the resulting parameter estimates (Appendix F gives the full, model-specific estimates).

For our preferred specification with model-specific initial MAC slopes and rates of TC, exogenous TC reduces the MAC slope by around 3% per year to a long-run value that is around 85% lower than its initial value. The slope halves every 25 years or so. For endogenous TC, the learning elasticity $\chi = 0.21$ in our preferred specification. This gives a learning rate $1 - 2^{-\chi} = 14\%$,

[13]The IPCC AR6 database is available at `https://data.ene.iiasa.ac.at/ar6/` and the NGFS Phase 5 transition database is available at `https://data.ene.iiasa.ac.at/ngfs/#/workspaces`.

[14]Alternatively, 45 different IAMs counting multiple members of the same family, e.g., different model versions, or energy models with and without coupling to land-use models.

[15]We group some of the 18 model families together to form 11 groups, to avoid non-convergence.

i.e., the percentage reduction in the MAC slope each time cumulative abatement doubles. Larger historical learning rates have been recorded for some technologies, such as photovoltaics (32%), wind power (19%) and battery technologies (42%), while lower learning rates have been recorded for other technologies, such as hydroelectricity and nuclear (both 0%) (Way et al., 2022).[16] Notice that for both types of TC, the fit on all models and the fit only on IAMs with TC of the same type give similar outcomes. For endogenous TC, the learning rate is faster but this is partially offset by lower cumulative emissions at time zero, A_0. Therefore, we find little evidence of a systematic relationship between the speed and the representation of TC.

4.2 Simple IAM

We calibrate and simulate a simple numerical IAM building on the theory models of Section 2, in particular the model with an aggregate abatement cost function. Welfare/utility is as in (1), however, to make the model more realistic we introduce population growth n at a decreasing rate, g_n. Warming is as in (2) and damages are as in (3). Emissions are as in (4), but redefined as greenhouse gas rather than just CO_2 emissions.[17] In addition, supported by the fact that BAU emissions scenarios in the IPCC/NGFS databases are roughly constant over this century, we assume that σ decreases at the same rate as output increases – AEEI cancels out the scale effect on emissions. Production is implicitly given by (6), but we assume the economy is approximately on a balanced growth path, so the economy grows at the rate of exogenous labour productivity. In the case of exogenous TC that reduces the slope of the MAC, the parameter φ falls according to $\varphi(t) = \varphi(\infty) + [\varphi(0) - \varphi(\infty)]\, e^{-g_\varphi t}$, i.e., the slope of the MAC function starts at $\varphi(0)$ and gradually approaches $\varphi(\infty)$ in the long run.

To avoid unrealistic jumps in initial emissions, we also add a penalty on the speed of abatement. This is a simple way to factor in the effect of adjustment costs and capital inertia, without explicitly modelling them. Rapid abatement may require costly repurposing/stranding of fossil-fuel-based capital, and green capital accumulation may also face bottlenecks. We assume a quadratic total speed penalty on final consumption, i.e., a linear marginal speed penalty, $\partial c(t)/\partial v(t) = \theta v c(t)$, where abatement speed is v and the abatement speed parameter θ is recovered from the structural estimation above (because quantitative/policy IAMs usually contain such constraints). Appendix C derives the effect of inertia on the model's optimality conditions.

The optimisation problem (for cost-benefit analysis) corresponds to:

$$\max_{\mu(t)} \int_0^\infty e^{-(\delta - n(t))t} \frac{c(t)^{1-\eta}}{1-\eta} dt, \tag{37}$$

$$\text{s.t.} \quad c(t) = \bar{c}(0) \exp\left[gt - \frac{\varphi(t)}{2} a(t)^2 \left(\frac{S_a(t)}{S_a(0)} \right)^{-\chi} - \frac{\theta}{2} v(t)^2 - \frac{\gamma}{2} \zeta^2 S(t)^2 \right], \tag{38}$$

[16] Our aggregate learning rate is not easy to compare with technology-specific learning rates, because not only learning rates χ_i but also initial levels of knowledge A_i are heterogenous. Our aggregate initial level of cumulative abatement A_0 assumes endogenous TC is slower in later decades (that is, green technology as a 'family' matures). Yet, there will be new fast-learning technologies in the future merely because they start with a low initial stock of knowledge $A_{0,i}$. Comparing the dynamics of an aggregate learning mechanism with disaggregated technologies as in Appendix B is a fruitful area of future research.

[17] We make a corresponding adjustment to the TCRE coefficient ζ.

	No TC	Exogenous TC				Endogenous, learning-based TC			
	All models	All models	Exog models	FEs: initial MACs	FEs: initial MACs & TC	All models	Endog models	FEs: initial MACs	FEs: initial MACs & TC
φ_0	2.7e-05***	1.2e-04***	1.2e-04***	1.1e-04***	1.1e-04***	1.4e-04***	1.3e-04***	9.1e-05***	1.1e-04***
g_φ		.048***	.049***	.013***	0.027***				
φ_∞		1.9e-05***	1.5e-05***	2.2e-07	1.7e-05***				
χ						.362***	.302***	.315***	0.211***
A_0						16.3***	12.7***	92.5***	36.6***
θ	1.5e-03***	-1.7e-05	-2.9e-05	7.7e-04***	4.7e-04***	5.1e-05	9.2e-05	6.6e-04***	4.5e-04***
N	12393	12393	6545	12393	12393	12393	5848	12393	12393
LL	4.9e+04	5e+04	2.7e+04	5.2e+04	5.2e+04	5.0e+04	2.3e+04	5.2e+04	5.2e+04

Table 2: Parameter estimates for fitting both total and marginal abatement costs to the IPCC and NGFS databases of IAM results. The model with exogenous TC has the following formula for the slope of the MAC function $\varphi_t = \varphi_\infty + (\varphi_0 - \varphi_\infty)\, e^{-g_\varphi t}$. In the fixed-effects (FEs) columns, we report mean coefficients and standard errors. *** indicates significance at the 1% level.

Parameter	Value	Source
δ	0.008	Drupp et al. (2018), 20% trimmed mean of experts
η	1.3	As above
n, $-g_n$	0.0105, 0.013	United Nations (2022)
g	0.02	By assumption
ζ	0.0006	IPCC AR6 WGI (IPCC, 2021)
γ	0.0154	Howard and Sterner (2017)
GDP_{2020}	US$84.537 trn	IMF (2021)
P_{BAU}	60 GtCO2eq	IPCC AR6 WGIII (IPCC, 2022)

Table 3: Exogenous parameter values for the numerical simulations.

where $\bar{c}$ is baseline consumption and $a(t) = Y(t)\sigma(t)\mu(t)$ denotes abatement. The model is solved as a boundary value problem with MATLAB's bvp5c function.[18] Table 3 reports the additional, exogenous parameter values that complement the TC parameter values recovered from the structural estimation. The model is run out to 2500 in order to avoid terminal values affecting the optimal paths during our period of interest.

4.3 Results

Cost-benefit analysis

Figure 1 plots optimal, cost-benefit climate policies for no TC (think of this as a straw-person scenario), exogenous TC and endogenous, learning-based TC. TC has a quantitatively large effect on optimal emissions, temperatures and MACs, and there is a substantial difference between the two forms of TC. The incentive to wait for abatement costs to fall is clear under exogenous TC. Emissions fall immediately, but they are higher than under no TC until 2036. Thereafter, they continue to fall in a roughly linear fashion and by 2100 they are much lower than under either of the other two scenarios.[19] Under endogenous, learning-based TC, emissions fall very rapidly in the first ten years (this despite the presence of capital inertia), illustrating the incentive to abate early in order to reduce future abatement costs. Thus, the different incentives of exogenous and endogenous, learning-based TC are evident in the comparison of their emissions profiles. Either form of TC results in temperatures at the end of the century that are substantially lower than under no TC – from 3.0°C under no TC to 2.5°C under exogenous TC and 2.4°C under endogenous, learning-based TC. Exogenous TC leads to a substantial reduction in the MAC/optimal carbon price, starting 26% lower than under no TC and with a widening gap over time. By contrast, under endogenous TC the optimal MAC starts 12% higher than under no TC but then grows more slowly, ending up 30% lower in 2100. Below we decompose this effect on the optimal MAC into the SCC (optimal carbon tax) and the marginal future learning gain (optimal deployment subsidy). The difference between the initial MAC under exogenous and endogenous, learning-based TC is \$73/tCO_2. Appendix G reports an equivalent set of results when capital inertia is ignored. Results are similar but large initial differences in emissions are optimal.

[18] For the cost-effectiveness analysis, we minimise total abatement costs choosing abatement every two years until 2200 using MATLAB's fmincon function

[19] Note that under exogenous TC, emissions are substantially negative in the very long run, beyond 2100.

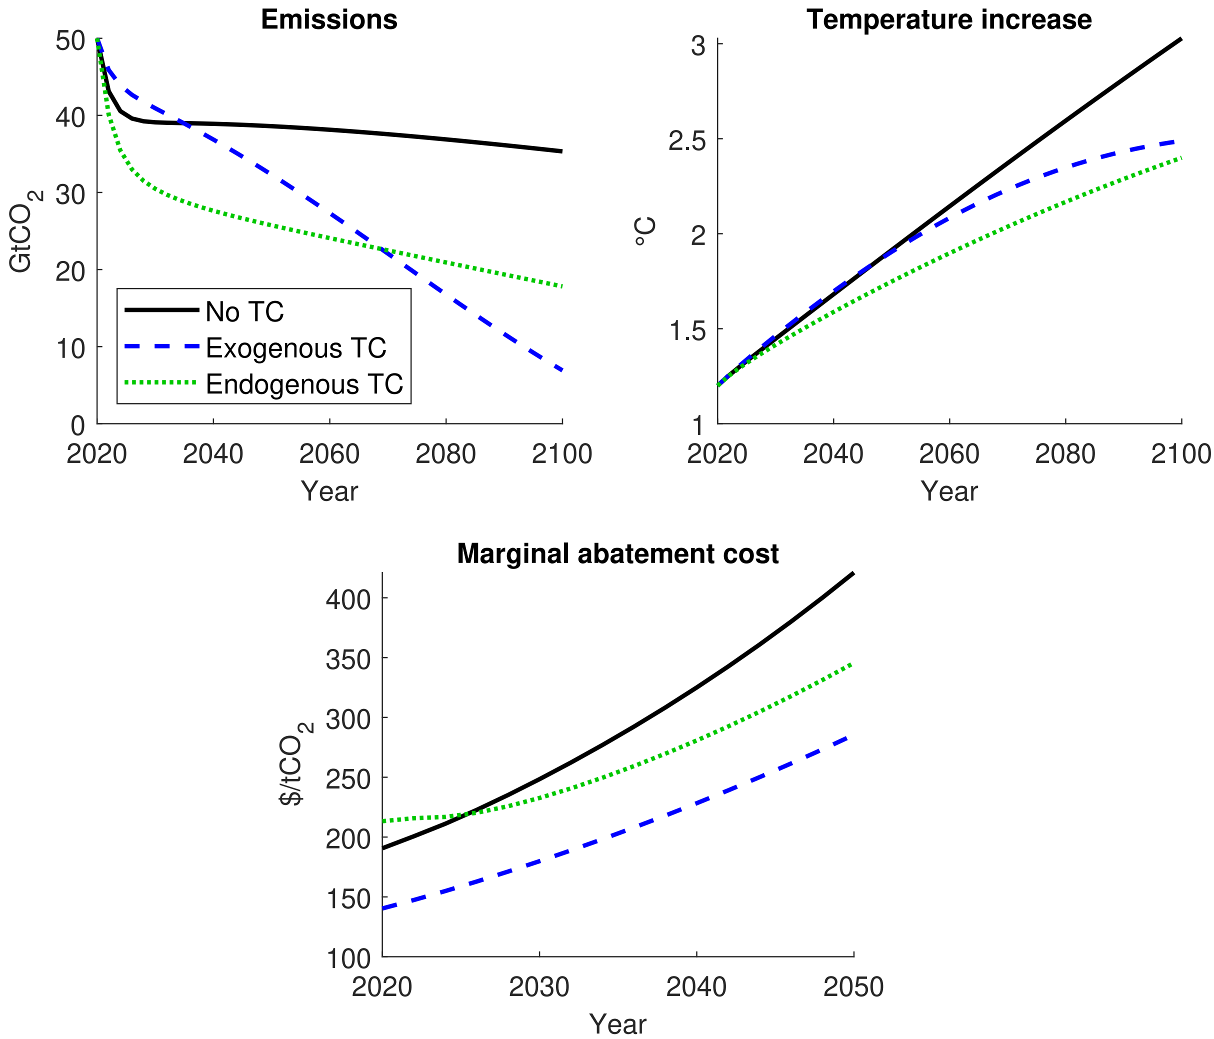

Figure 1: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

Cost-effectiveness analysis

Figure 2 plots optimal, cost-effective climate policies, imposing a temperature constraint of 1.75°C as representative of the UN Paris Agreement goal of "well below 2°C". This is a very demanding scenario in general, as evidenced by the rapidly falling emissions paths and high MACs. As in the cost-benefit case, exogenous TC leads to backloading abatement effort. Emissions are higher than under no TC until mid-century and lower thereafter. Consequently, temperatures hit the 1.75°C ceiling earlier. Compared to exogenous TC, endogenous, learning-based TC leads to more abatement in the next few decades and less abatement thereafter. Effort is marginally more frontloaded in order to realise endogenous future learning gains. However, unlike the cost-benefit case, endogenous, learning-based TC does not lead to lower emissions in the short run than under no TC, rather they are higher. This was identified as a possible outcome in the theory section. In a very ambitious abatement scenario, the cost-reduction effect of endogenous, learning-based TC outweighs the future learning effect. The difference in MACs due to TC is very large in this 1.75°C cost-effectiveness scenario. The MAC under exogenous TC is 26% lower

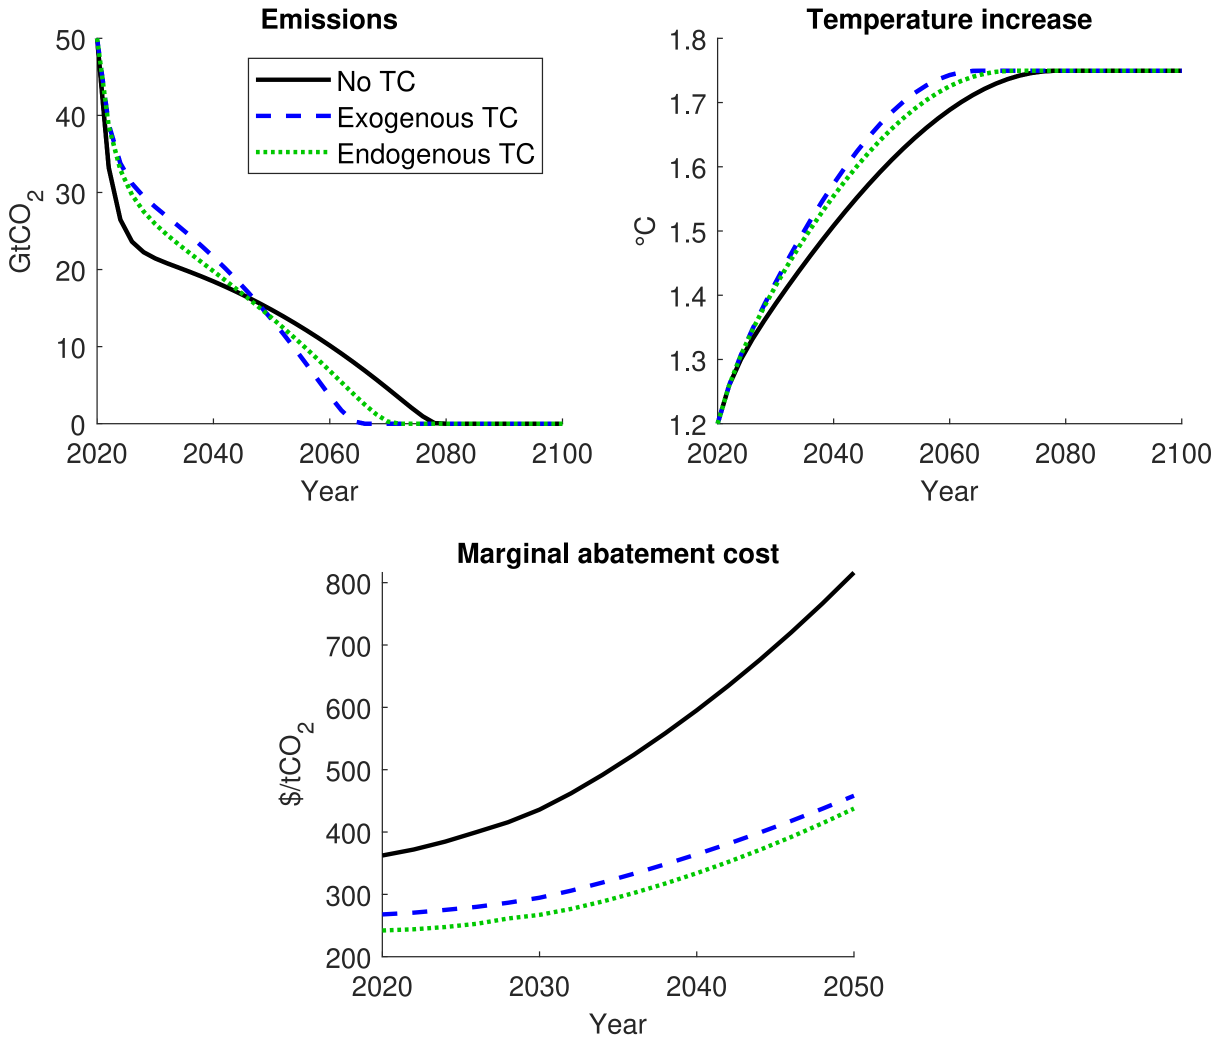

Figure 2: Optimal, cost-effective climate policies without TC, with exogenous TC and with endogenous TC. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

in 2020 than under no TC, widening to 44% in 2050. This is interpretable as the difference in the optimal carbon tax. The MAC under endogenous, learning-based TC is 33% lower in 2020 than under no TC and 46% lower in 2050. This contains both the optimal carbon tax and deployment subsidy.[20]

Estimating the endogenous future gain effect

The above figures show different effects of exogenous and endogenous TC on optimal climate policies. The main conceptual difference between the two TC processes is the endogenous future gain effect, i.e., internalising the marginal external gain in abatement knowledge. However, the above results do not isolate the endogenous future gain effect. So, Figure 3 decomposes the optimal MAC under endogenous, learning-based TC into its two components, the SCC/Hotelling price and the marginal future learning gain. An optimal policy would comprise a carbon tax

[20]Note that the MAC paths under the two forms of TC cross each other in 2060 (not shown in the figure), with the MAC under exogenous TC subsequently being lower than under endogenous, learning-based TC.

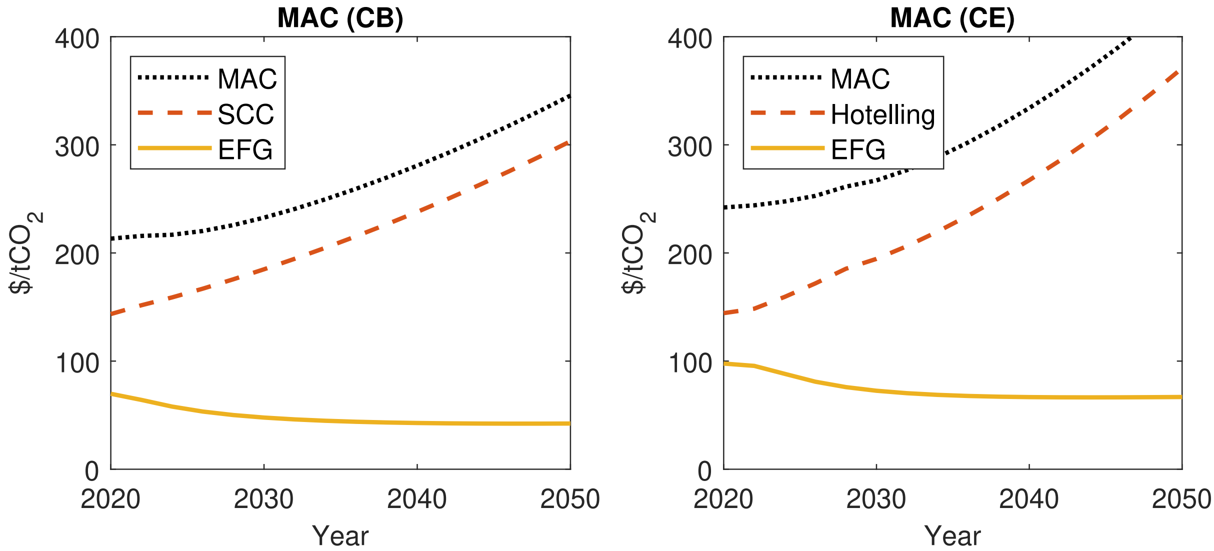

Figure 3: Decomposition of the optimal MAC into the SCC/Hotelling price and the endogenous future gain effect (EFG), for the cost-benefit (left) and cost-effectiveness (right) problems.

equal to the former and a deployment subsidy equal to the latter.

Looking first at the cost-benefit problem, the SCC is \$143/tCO_2 in 2020, \$185/tCO_2 in 2030 and \$303/tCO_2 in 2050. The marginal future learning gain is \$70/tCO_2 in 2020, falling to \$48/tCO_2 in 2030 and roughly flat thereafter. In the cost-effectiveness case, the Hotelling price is similar to the SCC, at \$144/tCO_2 in 2020, \$195/tCO_2 in 2030 and \$371/tCO_2 in 2050. But the marginal future learning gain is larger at \$98/tCO_2 in 2020 and \$73/tCO_2 in 2030. One might wonder why the initial Hotelling price is not much larger than the SCC given cumulative emissions are much lower. The answer is twofold: first, the Hotelling path increases at a steeper rate. This in turn increases the forward-looking endogenous learning gains, which make abatement cheaper.

But how does the dynamic incentive created by endogenous TC impact on optimal emissions, temperatures and MACs? This requires a different kind of analysis, because this dynamic incentive also affects the SCC. That is, take the endogenous future gain effect away and the SCC also changes. Here we develop a method of isolating the dynamic incentive of endogenous, learning-based TC. The method proceeds in two steps. The first step is to estimate a time-dependent MAC function (i.e., exogenous TC), with a MAC that is identical to the endogenous TC model at each point in time. This can be achieved using a sufficiently high-order polynomial. The second step is to take the exogenous-TC MAC curve so estimated, and use it to recalculate optimal model trajectories. Any difference between the optimal paths of the endogenous, learning-based TC model and its exogenous replica must then be down to the dynamic incentive of endogenous TC (see Appendix E.1).

Figure 4 plots the results. Emissions are 10% lower in 2050 and 5% lower in 2100 due to the endogenous future gain effect, leading to 0.13°C less warming at the end of the century. The MAC starts 41% higher and is 9% higher in 2050. These results indicate that, insofar as TC is endogenous and learning-based, models that do not include the incentive structure of a learning externality produce too little abatement, optimal MACs that are too low, and temperatures

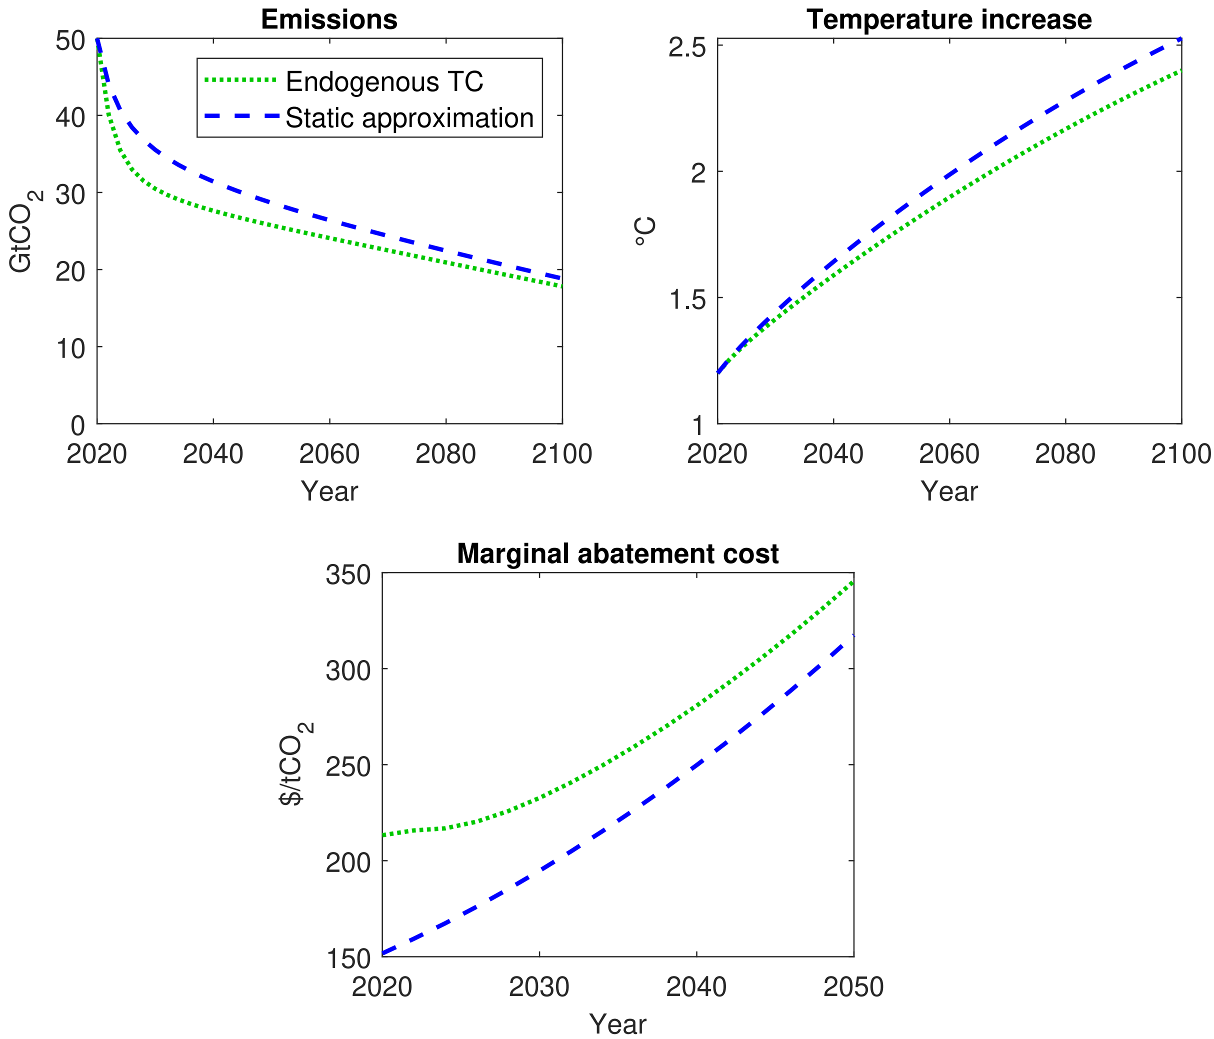

Figure 4: Comparison of a model of endogenous TC with a model without TC, but where the static MAC function is calibrated to give the same marginal cost and quantity of abatement at each point in time as the endogenous TC model. The difference between the paths is exactly the endogenous future gain effect. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

that are too high. In our review of models, we showed that TC is exogenous in most IAMs, but this observation also applies to models of semi-endogenous TC, as these also omit the dynamic incentive. The estimates here provide an upper bound on the bias though, because they are derived from comparing a model of pure exogenous TC and a model of pure endogenous TC. Reality likely incorporates both.

Appendix E.2 further compares a model of exogenous TC with a static model with identical MACs at each point in time. For example, if, say, the exogenous TC model projects zero emissions in 2050 at a marginal cost of \$250/tCO_2, the static model would have the same MAC at zero emissions. The static MAC function will be concave, since TC makes the linear instantaneous MAC function fall over time. We show theoretically that the static approximation of the exogenous TC model is exact under certain assumptions and quantitatively that it is almost exact under more general assumptions. By contrast, a static model cannot imitate the dynamics of a model with endogenous, learning-based TC.

Sensitivity analysis

Here we analyse the sensitivity of the optimal, cost-benefit policy to variation in the abatement cost/TC parameters, and the discount rate. For the former variation, we use the standard errors from the GMM estimates in Table 2. For the latter, we use the range of responses from the expert survey on discounting by Drupp et al. (2018). The results are summarised in Table 4, with detailed results contained in Appendix G. There are four scenarios: low discount rate; high discount rate; slow TC; fast TC.

Reducing the discount rate leads to lower emissions/temperatures and higher MACs, as expected. Increasing the discount rate leads to the opposite effects. The effects of TC are qualitatively the same under different discount rates, except that with a low discount rate, endogenous, learning-based TC leads to a lower initial MAC than under no TC, not a higher one. The explanation for this is the same as in the cost-effectiveness case above: when emissions fall very fast, the cost reduction effect outweighs the endogenous learning effect.

If TC is slow, emissions/temperatures are higher, but so are MACs.[21] The opposite is true for fast TC. The comparative effects of exogenous and endogenous TC are qualitatively the same whether TC is slow or fast.[22]

5 Discussion

Our aim in this paper has been to assess, both qualitatively and quantitatively, the effect of different representations of TC in the IAM literature on optimal climate policy. In doing so, we have surveyed a wide range of models distributed across several strands of literature, including mostly theoretical contributions to environmental economics and applied, quantitative contributions to policy. TC has been represented in multiple ways and has complex dynamic effects on emissions, temperatures, costs, taxes and subsidies. We have tried to make sense of the different approaches by setting up a common theoretical and quantitative framework, within which different TC mechanisms can be analysed.

TC matters. Exogenous TC that makes future abatement cheaper creates an incentive to abate less in the short run and more in the long run if temperature is optimally chosen. Emissions are much lower in the long run according to our quantitative analysis, resulting in 0.5°C less warming in 2100. If temperature is constrained, exogenous TC that makes future abatement cheaper creates a similar incentive to backload abatement and hit the temperature ceiling faster. If exogenous TC happens instead through autonomous decarbonisation, it plausibly reduces emissions in the short run but increases them in the long run. Again, if temperature is constrained, abatement is backloaded and the temperature ceiling is hit faster. Endogenous TC based on learning-by-doing also makes future abatement cheaper, but there is the additional and opposing learning effect, which incentivises early abatement. Our quantitative results suggests that under cost-benefit analysis, the latter effect outweighs the former so that emissions and temperatures are always lower than under no TC. The results also suggest that the optimal

[21]Under endogenous, learning-based TC, the MAC is lower under low TC for just the first two years.

[22]Care should be taken in comparing exogenous and endogenous, learning-based TC under low/high rates, however – setting the parameters at plus/minus one standard deviation is not assured to produce comparable rates of TC.

	Emissions ($GtCO_2$)			Temperatures (°C)			MACs ($/$tCO_2$)	
	2030	2050	2100	2030	2050	2100	2020	2050
	Main specification							
No TC	39.10	38.60	35.33	1.45	1.91	3.03	190.62	421.05
Exogenous TC	40.93	32.31	6.91	1.46	1.90	2.49	140.32	285.47
Endogenous TC	30.48	25.74	17.81	1.41	1.75	2.40	213.26	345.54
	Low discount rate							
No TC	17.80	20.03	19.65	1.34	1.57	2.18	383.72	745.72
Exogenous TC	30.81	21.11	-3.79	1.42	1.73	1.97	224.63	394.30
Endogenous TC	9.89	6.59	4.34	1.32	1.41	1.58	343.94	472.22
	High discount rate							
No TC	50.00	48.90	45.57	1.50	2.09	3.51	86.22	221.19
Exogenous TC	48.88	42.26	19.64	1.50	2.05	2.96	76.62	184.71
Endogenous TC	45.37	41.41	32.56	1.48	2.00	3.11	107.23	216.68
	Slow technical change							
No TC	39.10	38.60	35.33	1.45	1.91	3.03	190.62	421.05
Exogenous TC	40.94	33.35	9.92	1.46	1.91	2.56	143.75	294.66
Endogenous TC	32.65	29.20	22.46	1.42	1.79	2.57	209.88	367.71
	Fast technical change							
No TC	39.10	38.60	35.33	1.45	1.91	3.03	190.62	421.05
Exogenous TC	40.88	31.21	4.29	1.46	1.90	2.42	137.29	277.34
Endogenous TC	28.27	22.03	12.96	1.41	1.70	2.22	215.44	321.47

Table 4: Sensitivity analysis on the discount rate and the rate of technical change. The low discount rate corresponds to $\delta = 0$ and the high discount rate to $\delta = 0.02$. Further sensitivity analysis on η can be found in Appendix G. Fast (slow) TC corresponds to plus (minus) one standard deviation for g_φ in the model of exogenous TC and for χ and A_0 in the model of endogenous, learning-based TC.

paths under exogenous, MAC-reducing TC and endogenous, learning-based TC diverge a lot. Under cost-benefit analysis, emissions, temperatures and MACs are substantially different. Under cost-effectiveness analysis with a tight temperature constraint, emissions and temperatures are necessarily more similar, but MACs are even more dissimilar. We decompose the difference in the MAC into the difference in the SCC and the difference in the learning externality, with implications for optimal Pigouvian taxes and learning subsidies respectively. Models of endogenous, R&D-based TC create an incentive for early action through subsidising R&D, but optimal emissions, temperature and MAC/carbon tax paths are plausibly isomorphic with exogenous TC. The two forms of endogenous TC that we consider, based on learning-by-doing or R&D, can imply different technology cost dynamics, an issue we feel deserves further analysis and discussion in the future. We try to get this discussion started with an analytical comparison of technology cost paths and an illustration of how learning and R&D can be combined within a single model of relatively low complexity.

Sensitivity analysis reveals that our results are qualitatively robust to variation in the rate of TC (and the discount rate). The sensitivity analysis on the rate of TC, carried out using the standard errors of our statistical estimates, also provides an insight into the likely impacts of TC on optimal trajectories, in case the IAM literature as a whole underestimates future technology cost reductions. This is an implication of recent work by Way et al. (2022), who argue IAMs underestimate deployment rates for renewable energy technologies and overestimate their costs. Supposing this critique has some validity, we might then be more guided by our results for high-end TC as calibrated on the IAM databases.[23] Mechanically, fast TC, so defined and measured, results in a larger quantitative effect on optimal emissions, temperatures and carbon prices.

Uncertainty about TC is an important feature of the problem. Although there exist many papers with sensitivity analysis on technological parameters, like ours, to the best of our knowledge the literature has not yet investigated optimal emissions under uncertainty in a dynamic model where information on TC is gradually discovered. A fully dynamic stochastic model goes beyond the scope of this paper, but from our analytical solutions we can speculate about the effect of uncertainty.

Since total abatement costs are convex in abatement, Jensen's inequality indicates that the expected value of future abatement costs increases under uncertainty. This increases future abatement costs relative to current, certain abatement costs, and is an argument for earlier abatement with a higher initial MAC. The higher MAC can be understood as a risk premium society is willing to pay to insure against the possible outcome of slower-than-expected TC.

But how much this matters depends as usual on whether TC is exogenous or endogenous, and whether the objective is cost-benefit or cost-effectiveness. When TC is exogenous and the objective is cost-effectiveness, the above effect is the only one at play and in principle the risk premium on TC uncertainty could be quantitatively large. In the case of endogenous TC, we would expect the risk premium to be lower, because we know from our analysis that the

[23]In our fast TC scenario, abatement costs are reduced by one third between 2020 and 2040. This corresponds to the cost reduction for wind technologies in the fast transition scenario of Way et al. (2022). Their cost reductions for batteries and electrolysers are even larger because these technologies start with a very low cumulative installed capacity. By contrast, they argue that other technologies such as carbon capture and storage and nuclear have seen almost no cost reductions over the last decades.

endogenous future gain effect attenuates the effect of TC uncertainty on the initial MAC. In other words, since endogenous TC has a smaller impact on the optimal MAC (compared to exogenous TC), the societal cost of wrongly anticipating TC is lower. Similarly, in a cost-benefit analysis, lower-than-expected TC leads to a higher optimal peak temperature, attenuating the initial price/MAC adjustment. Again, since the effect of TC on the initial MAC is lower than under cost-effectiveness, the social cost of wrongly anticipating TC and the corresponding risk premium are lower. Combining both effects, i.e. in a cost-benefit analysis with endogenous TC, the effect of TC on the initial MAC should be small.

Note that these are mere first-order effects of adding uncertainty on parameters χ and g_φ. There are other effects that go beyond this intuition, for example when the uncertainty regarding the growth rate of the economy is correlated with uncertainty regarding TC. Also, uncertainty affects investment incentives. In the case of long-term, irreversible investments with large uncertainty over the profitability of the technology, uncertainty leads to an incentive to postpone the investment and wait for new information, a.k.a. option value, the value of keeping options open. For example, uncertainty about the availability of nuclear fusion in the future can create an incentive to postpone the investment in a new nuclear fission plant today. Investigating optimal emissions under technological uncertainty is therefore a fruitful avenue for future research.

References

Acemoglu, D., Aghion, P., Bursztyn, L., Hemous, D., 2012. The environment and directed technical change. American Economic Review 102, 131–66.

Acemoglu, D., Akcigit, U., Hanley, D., Kerr, W., 2016. Transition to clean technology. Journal of Political Economy 124, 52–104.

Bibas, R., Cassen, C., Crassous, R., Guivarch, C., Hamdi-Cherif, M., Hourcade, J.C., Leblanc, F., Méjean, A., Broin, E.Ó., Rozenberg, J., et al., 2022. IMpact Assessment of CLIMate policies with IMACLIM-R 1.1. Model documentation version 1.1.

van den Bijgaart, I., Gerlagh, R., Liski, M., 2016. A simple formula for the social cost of carbon. Journal of Environmental Economics and Management 77, 75–94.

Böhringer, C., Löschel, A., Moslener, U., Rutherford, T.F., 2009a. EU climate policy up to 2020: an economic impact assessment. Energy Economics 31, S295–S305.

Böhringer, C., Löschel, A., Rutherford, T.F., 2009b. Policy analysis based on computable equilibrium (PACE), in: Bosetti, V., Gerlagh, R., Schleicher, S.P. (Eds.), Modelling Sustainable Development. Edward Elgar Publishing.

Böhringer, C., Peterson, S., Rutherford, T.F., Schneider, J., Winkler, M., 2021. Climate policies after paris: Pledge, trade and recycle: Insights from the 36th energy modeling forum study (emf36). Energy Economics 103, 105471.

Bond-Lamberty, B., Patel, P., Lurz, J., pkyle, kvcalvin, Smith, S., abigailsnyder, Dorheim, K.R., russellhz, mbins, Link, R., skim301, nealtg, Narayan, K., Turner, S.W.D., S., A., Feng, L., enlochner, cwroney, Lynch, C., jhoring, Khan, Z., siddarthd96, orourkepr, JonathanHuster, Haewon, Ou, Y., Iyer, G., mwisepnnl, marideeweber, 2022. Jgcri/gcam-core: Gcam 6.0. doi:10.5281/zenodo.6619287.

Bramoullé, Y., Olson, L.J., 2005. Allocation of pollution abatement under learning by doing. Journal of Public Economics 89, 1935–1960.

Bretschger, L., Ramer, R., Schwark, F., 2011. Growth effects of carbon policies: applying a fully dynamic cge model with heterogeneous capital. Resource and Energy Economics 33, 963–980.

Buonanno, P., Carraro, C., Galeotti, M., 2003. Endogenous induced technical change and the costs of kyoto. Resource and Energy economics 25, 11–34.

Burniaux, J.M., Truong, T.P., 2002. GTAP-E: an energy-environmental version of the GTAP model. Technical Report. GTAP Technical Papers.

Calvin, K., Bond-Lamberty, B., Clarke, L., Edmonds, J., Eom, J., Hartin, C., Kim, S., Kyle, P., Link, R., Moss, R., McJeon, H., Patel, P., Smith, S., Waldhoff, S., Wise, M., 2017. The SSP4: a world of deepening inequality. Global Environmental Change 42, 284–296.

Cambridge Econometrics, 2019. E3ME technical manual v6.1.

Capros, P., Van Regemorter, D., Paroussos, L., Karkatsoulis, P., Fragkiadakis, C., Tsani, S., Charalampidis, I., Revesz, T., Perry, M., Abrell, J., 2013. GEM-E3 model documentation. JRC Scientific and Policy Reports 26034.

Château, J., Dellink, R., Lanzi, E., 2014. An overview of the OECD ENV-linkages model. Technical Report. OECD.

Chen, Y.H., Paltsev, S., Reilly, J.M., Morris, J.F., Babiker, M.H., 2015. The MIT EPPA6 model: economic growth, energy use, and food consumption. Technical Report. MIT Joint Program on the Science and Policy of Global Change.

Clarke, L., Jiang, K., Akimoto, K., et al., 2014. Asessing transformation pathways, in: Edenhofer, O., Pichs-Madruga, R., Sokona, Y., et al. (Eds.), Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA.

Collins, M., Knutti, R., Arblaster, J., Dufresne, J.L., Fichefet, T., Friedlingstein, P., Gao, X., Gutowski, W., Johns, T., Krinner, G., Shongwe, M., Tebaldi, C., A.J., W., Wehner, M., 2013. Long-term climate change: projections, commitments and irreversibility, in: Stocker, T., Qin, C., Plattner, G.K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgeley, P. (Eds.), Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA.

Corong, E.L., Hertel, T.W., McDougall, R., Tsigas, M.E., Van Der Mensbrugghe, D., 2017. The standard GTAP model, version 7. Journal of Global Economic Analysis 2, 1–119.

Di Maria, C., Smulders, S., 2017. A paler shade of green: Environmental policy under induced technical change. European economic review 99, 151–169.

Dietz, S., van der Ploeg, F., Rezai, A., Venmans, F., 2021. Are economists getting climate dynamics right and does it matter? Journal of the Association of Environmental and Resource Economists 8, 895–921.

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.

Drupp, M.A., Freeman, M.C., Groom, B., Nesje, F., 2018. Discounting disentangled. American Economic Journal: Economic Policy 10, 109–34.

Eboli, F., Parrado, R., Roson, R., 2010. Climate-change feedback on economic growth: explorations with a dynamic general equilibrium model. Environment and Development Economics 15, 515–533.

Emmerling, J., Drouet, L., Reis, L.A., Bevione, M., Berger, L., Bosetti, V., Carrara, S., De Cian, E., D'Aertrycke, G.D.M., Longden, T., et al., 2016. The WITCH 2016 model-documentation and implementation of the shared socioeconomic pathways. Technical Report.

Fricko, O., Havlik, P., Rogelj, J., Klimont, Z., Gusti, M., Johnson, N., Kolp, P., Strubegger, M., Valin, H., Amann, M., Ermolieva, T., Forsell, N., Herrero, M., Heyes, C., Kindermann, G., Krey, V., McCollum, D.L., Obersteiner, M., Pachauri, S., Rao, S., Schmid, E., Schoepp, W., Riahi, K., 2017. The marker quantification of the Shared Socioeconomic Pathway 2: a middle-of-the-road scenario for the 21st century. Global Environmental Change 42, 251–267.

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

Fujimori, S., Hasegawa, T., Masui, T., 2017a. Aim/cge v2.0: Basic feature of the model, in: Fujimori, S., Kainuma, M., Masui, T. (Eds.), Post-2020 Climate Action: Global and Asian Perspectives. Springer Singapore, Singapore, pp. 305–328.

Fujimori, S., Hasegawa, T., Masui, T., Takahashi, K., Herran, D.S., Dai, H., Hijioka, Y., Kainuma, M., 2017b. SSP3: AIM implementation of Shared Socioeconomic Pathways. Global Environmental Change 42, 268–283.

Gavard, C., Voigt, S., Genty, A., 2022. Using emissions trading schemes to reduce heterogeneous distortionary taxes: the case of recycling carbon auction revenues to support renewable energy. Energy Policy 168, 113–133.

Gerlagh, R., Kverndokk, S., Rosendahl, K.E., 2009. Optimal timing of climate change policy: interaction between carbon taxes and innovation externalities. Environmental and Resource Economics 43, 369–390.

Gillingham, K., Newell, R.G., Pizer, W.A., 2008. Modeling endogenous technological change for climate policy analysis. Energy Economics 30, 2734–2753.

Gillingham, K., Nordhaus, W., Anthoff, D., Blanford, G., Bosetti, V., Christensen, P., McJeon, H., Reilly, J., 2018. Modeling uncertainty in integrated assessment of climate change: a multimodel comparison. Journal of the Association of Environmental and Resource Economists 5, 791–826.

Goettle, R.J., Ho, M.S., Jorgenson, D.W., Slesnick, D.T., Wilcoxen, P.J., EP-W, E.C., 2007. IGEM, an inter-temporal general equilibrium model of the US economy with emphasis on growth, energy and the environment. Technical Report Model documentation in support of ongoing policy analysis for the US Environmental Protection Agency, EPA Contract No. EP-W-05-035.

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica 82, 41–88.

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.

Greaker, M., Pade, L.L., 2009. Optimal carbon dioxide abatement and technological change: should emission taxes start high in order to spur R&D? Climatic Change 96, 335–355.

Grubb, M., Köhler, J., Anderson, D., 2002. Induced technical change in energy and environmental modeling: analytic approaches and policy implications. Annual Review of Energy and the Environment 27, 271–308.

Grubb, M., Lange, R.J., Cerkez, N., Sognnaes, I., Wieners, C., Salas, P., 2024. Dynamic determinants of optimal global climate policy. Structural Change and Economic Dynamics 71, 490–508.

Grubb, M., Wieners, C., Yang, P., 2021. Modeling myths: on DICE and dynamic realism in integrated assessment models of climate change mitigation. Wiley Interdisciplinary Reviews: Climate Change 12, e698.

Hart, R., 2019. To everything there is a season: Carbon pricing, research subsidies, and the transition to fossil-free energy. Journal of the Association of Environmental and Resource Economists 6, 349–389.

Hassler, J., Krusell, P., Olovsson, C., 2021. Directed technical change as a response to natural resource scarcity. Journal of Political Economy 129, 3039–3072.

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.

IEA, 2021. World Energy Model Documentation.

IMF, 2021. World Economic Outlook Database.

IPCC, 2018. Global Warming of 1.5C: An IPCC Special Report on the Impacts of Global Warming of 1.5C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Cambridge University Press, Cambridge, UK and New York, NY, USA. URL: https://www.ipcc.ch/sr15/.

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. URL: https://report.ipcc.ch/ar6/wg1/IPCC_AR6_WGI_FullReport.pdf, doi:10.1017/9781009157896.

IPCC, 2022. Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA. URL: https://www.ipcc.ch/report/ar6/wg3/downloads/report/IPCC_AR6_WGIII_FullReport.pdf, doi:10.1017/9781009157926.

Jacoby, H.D., Reilly, J.M., McFarland, J.R., Paltsev, S., 2006. Technology and technical change in the MIT EPPA model. Energy Economics 28, 610–631.

Jakeman, G., Hanslow, K., Hinchy, M., Fisher, B.S., Woffenden, K., 2004. Induced innovations and climate change policy. Energy Economics 26, 937–960.

Jorgenson, D.W., Wilcoxen, P.J., 1993. Reducing us carbon emissions: an econometric general equilibrium assessment. Resource and Energy Economics 15, 7–25.

Keramidas, K., Kitous, A., Després, J., Schmitz, A., Diaz Vazquez, A., Mima, S., et al., 2017. POLES-JRC model documentation. Publications Office of the European Union 10, 225347.

Krey, V., Havlik, P., Kishimoto, P., Fricko, O., Zilliacus, J., Gidden, M., Strubegger, M., Kartasasmita, G., Ermolieva, T., Forsell, N., et al., 2020. MESSAGEix-GLOBIOM Documentation-2020 release. Technical Report. IIASA.

Kuik, O., Brander, L., Tol, R.S., 2009. Marginal abatement costs of greenhouse gas emissions: a meta-analysis. Energy Policy 37, 1395–1403.

Löschel, A., 2002. Technological change in economic models of environmental policy: a survey. Ecological Economics 43, 105–126.

Loulou, R., Goldstein, G., Kanudia, A., Lehtila, A., Remne, U., 2016. Documentation for the TIMES model-Part 1 2016. Technical Report.

Lucena, A.F., Hejazi, M., Vasquez-Arroyo, E., Turner, S., Köberle, A.C., Daenzer, K., Rochedo, P.R., Kober, T., Cai, Y., Beach, R.H., et al., 2018. Interactions between climate change mitigation and adaptation: the case of hydropower in Brazil. Energy 164, 1161–1177.

Luderer, G., Leimbach, M., Bauer, N., Kriegler, E., Baumstark, L., Bertram, C., Giannousakis, A., Hilaire, J., Klein, D., Levesque, A., et al., 2015. Description of the REMIND model (Version 1.6). Technical Report.

Manne, A., Richels, R., 2004. The impact of learning-by-doing on the timing and costs of CO2 abatement. Energy Economics 26, 603–619.

Mercure, J.F., Pollitt, H., Edwards, N.R., Holden, P.B., Chewpreecha, U., Salas, P., Lam, A., Knobloch, F., Vinuales, J.E., 2018. Environmental impact assessment for climate change policy with the simulation-based integrated assessment model E3ME-FTT-GENIE. Energy Strategy Reviews 20, 195–208.

Messner, S., 1997. Endogenized technological learning in an energy systems model. Journal of Evolutionary Economics 7, 291–313.

Nikas, A., Doukas, H., Papandreou, A., 2019. A detailed overview and consistent classification of climate-economy models, in: Doukas, H., Flamos, A., Lieu, J. (Eds.), Understanding Risks and Uncertainties in Energy and Climate Policy: Multidisciplinary Methods and Tools for a Low Carbon Society. Springer.

Nordhaus, W.D., 2017. Revisiting the social cost of carbon. Proceedings of the National Academy of Sciences 114, 1518–1523.

Octaviano, C., Paltsev, S., Gurgel, A.C., 2016. Climate change policy in Brazil and Mexico: Results from the MIT EPPA model. Energy Economics 56, 600–614.

Pant, H., 2007. Global Trade and Environment Model (GTEM). Technical Report. Australian Bureau of Agricultural and Resource Economics and Sciences Canberra, Australia.

Parrado, R., De Cian, E., 2014. Technology spillovers embodied in international trade: Intertemporal, regional and sectoral effects in a global CGE framework. Energy Economics 41, 76–89.

Popp, D., 2004. ENTICE: endogenous technological change in the DICE model of global warming. Journal of Environmental Economics and Management 48, 742–768.

Rezai, A., Van der Ploeg, F., 2016. Intergenerational inequality aversion, growth, and the role of damages: Occam's rule for the global carbon tax. Journal of the Association of Environmental and Resource Economists 3, 493–522.

Riahi, K., Schaeffer, R., Arango, J., Calvin, K., Guivarch, C., Hasegawa, T., Jiang, K., Kriegler, E., Matthews, R., Peters, G., Rao, A., Robertson, S., Sebbit, A., Steinberger, J., Tavoni, M., Van Vuuren, D., 2022. Mitigation pathways compatible with long-term goals, in: Shukla, P., Skea, J., Slade, R., Khourdajie, A.A., van Diemen, R., McCollum, D., Pathak, M., Some, S., Vyas, P., Fradera, R., Belkacemi, M., Hasija, A., Lisboa, G., Luz, S., Malley, J. (Eds.), Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA. book section 3. URL: https://www.ipcc.ch/report/ar6/wg3/downloads/report/IPCC_AR6_WGIII_Chapter03.pdf, doi:10.1017/9781009157926.005.

Sano, F., Akimoto, K., Homma, T., Tomoda, T., 2006. Analysis of technological portfolios for co2 stabilizations and effects of technological changes. The Energy Journal 27, 141–162.

Smulders, S., De Nooij, M., 2003. The impact of energy conservation on technology and economic growth. Resource and Energy Economics 25, 59–79.

Stehfest, E., van Vuuren, D., Bouwman, L., Kram, T., et al., 2014. Integrated assessment of global environmental change with IMAGE 3.0: model description and policy applications. Technical Report. Netherlands Environmental Assessment Agency (PBL).

Sue Wing, I., 2006. Representing induced technological change in models for climate policy analysis. Energy Economics 28, 539–562.

Sue Wing, I., Daenzer, K., Fisher-Vanden, K., Calvin, K., 2011. Phoenix model documentation. Technical Report. Maryland: Joint Global Change Research Institute, Pacific Northwest National Laboratory.

Traeger, C.P., 2023. Ace – analytic climate economy. American Economic Journal: Economic Policy 15, 372–406.

United Nations, 2022. World Population Prospects 2022, Online Edition. United Nations, Department of Economic and Social Affairs, Population Division.

van Vuuren, D.P., van der Wijst, K.I., Marsman, S., van den Berg, M., Hof, A.F., Jones, C.D., 2020. The costs of achieving climate targets and the sources of uncertainty. Nature Climate Change 10, 329–334.

Wada, K., Sano, F., Akimoto, K., Homma, T., 2012. Assessment of Copenhagen pledges with long-term implications. Energy Economics 34, S481–S486.

Way, R., Ives, M.C., Mealy, P., Farmer, J.D., 2022. Empirically grounded technology forecasts and the energy transition. Joule 6, 2057–2082.

Weyant, J., 2017. Some contributions of integrated assessment models of global climate change. Review of Environmental Economics and Policy 11, 115–137.

Wilkerson, J.T., Cullenward, D., Davidian, D., Weyant, J.P., 2013. End use technology choice in the national energy modeling system (nems): An analysis of the residential and commercial building sectors. Energy Economics 40, 773–784.

Van der Zwaan, B.C., Gerlagh, R., Schrattenholzer, L., et al., 2002. Endogenous technological change in climate change modelling. Energy Economics 24, 1–19.

Appendices for Online Publication

In Appendices A-E, we formally derive the analytical results in Section 2. Our strategy is to use the most general possible model. That is, to the extent possible we derive results for general functional forms, which then apply to the specific functions used in Section 2. This tends to keep the mathematical notation uncluttered and provides a set of results of potentially broad use, while maintaining a model structure in the main text that is more immediately relatable to the literature. Some results hold under more general conditions than others, therefore, as we proceed through the analysis we will make more specialised assumptions. Appendices F and G contain further results from our quantitative analysis.

Appendix A General model

For ease of notation, we omit time subscripts unless confusion is possible. Other subscripts are used to indicate partial derivatives. Therefore, we rename cumulative abatement A, instead of S_a in the main text.

Define a knowledge stock H that accumulates according to the following equation of motion, which depends on time t, the existing knowledge stock, investment I, and abatement $a = P_{BAU} - P$:

$$\dot{H} = \psi(t, H, I, a). \tag{39}$$

The most intuitive interpretation of H is a knowledge stock, but it can be any technological parameter that is path-dependent. Assume the function ψ is twice differentiable in all its arguments. The presence of H in the function may represent standing on the shoulders of giants ($\psi_H > 0$) or fishing out ($\psi_H < 0$). We assume $\psi_H < r$, i.e., the consumption discount rate, to avoid a bang-bang solution. The knowledge stock may increase over time ($\frac{\partial \psi}{\partial t} \geq 0$), e.g., via technological spillovers from non-green sectors, as a response to green R&D investment ($\psi_I \geq 0; \psi_{II} \leq 0$), or from learning-by-doing as a result of clean technology deployment ($\psi_a \geq 0; \psi_{aa} < 0$).

As in the main text, we assume that warming is proportional to cumulative emissions, that is, $T = \zeta S$ with

$$S_t = S_0 + \int_0^t P_u du \Leftrightarrow \dot{S} = P. \tag{40}$$

A.1 Cost-benefit analysis

Consider a consumption function $c\,(a, H, I, T, t)$, twice differentiable in all its arguments, where positive abatement is costly $-c_a \geq 0$, the MAC function is increasing ($-c_{aa} > 0$), and emissions beyond BAU are useless $-c_a|_{a \leq 0} = 0$. Knowledge decreases total and marginal abatement costs ($c_H > 0; -c_{aH} > 0$), and green R&D investment reduces consumption ($c_I = 1$). Climate warming causes convex damages ($c_T < 0, c_{TT} < 0$). We assume that apart from abatement costs, green R&D and damages, the economy is on a balanced growth path with a constant savings rate as in Dietz and Venmans (2019). As a result, we do not distinguish between production and consumption.

Population at time zero is normalised to one and grows at rate n. The utility function has the standard properties $u_c > 0, u_{cc} < 0$. The social planner maximises welfare as discounted utility,

$$\max_{\{a,I\}} \int_0^\infty e^{-(\delta-n)t} u\left(c\left(a, H, I, T, t\right)\right) dt, \tag{41}$$

subject to

$$\dot{S} = P_{BAU} - a; \dot{H} = \psi(t, H, I, a); S_0, H_0 \text{ given.} \tag{42}$$

The current value Hamiltonian of the problem is

$$\mathcal{H} = u\left(c\left(a, H, I, T(S), t\right)\right) - \lambda\left(P_{BAU} - a\right) + \mu\psi\left(t, H, I, a\right). \tag{43}$$

The FOCs include

$$u_c c_a + \mu\psi_a = \lambda, \tag{44}$$

$$\dot{\lambda} = \left(\delta - n\right)\lambda - u_c \zeta c_T, \tag{45}$$

$$u_c = \mu\psi_I, \tag{46}$$

$$\dot{\mu} = \left(\delta - n\right)\mu - u_c c_H - \mu\psi_H. \tag{47}$$

The transversality conditions are $\lim_{t\to\infty}\lambda e^{-(\delta-n)t} = 0$ and $\lim_{t\to\infty}\mu e^{-(\delta-n)t} = 0$.

Integrating (45) and (47), and plugging the integrals into (44), we obtain

$$c_{a_t} = \int_t^\infty e^{-(\delta-n)(u-t)} \frac{u_{c_u}}{u_{c_t}} \zeta\left(-c_{T_u}\right) du + \psi_{a_t} \int_t^\infty e^{-\int_t^u ((\delta-n)+\psi_H) ds} \frac{u_{c_u}}{u_{c_t}} c_{H_u} du. \tag{48}$$

Acknowledging that $\ln\frac{u_{c_u}}{u_{c_t}} = \int_t^u \frac{\dot{u_c}}{u_c} ds$ and defining the discount rate $r = \delta - n - \frac{\dot{u_c}}{u_c}$, we obtain that the MAC equals the present value of future marginal damages (the SCC) plus the future learning gain from an extra tonne of abatement today:

$$\underbrace{c_{a_t}}_{MAC} = \underbrace{\int_t^\infty e^{-\int_t^u r ds} \zeta\left(-c_{T_u}\right) du}_{SCC} + \underbrace{\overbrace{\psi_{a_t}}^{knowledge\ increment} \overbrace{\int_t^\infty e^{-\int_t^u (r+\psi_H) ds} c_{H_u} du}^{...and\ its\ effect\ on\ abatement\ costs}}_{Endog.\ Future\ Gain}. \tag{49}$$

This corresponds to Equation (12) in the main text, i.e., the optimality condition for the model with endogenous TC based on learning-by-doing. In the absence of learning-by-doing, the endogenous future gain term drops out, giving the optimality condition for exogenous TC in Equation (9).

Differentiating (44) and combining it with the other FOC gives

$$\dot{c_a} = r c_a + \frac{\psi_H \psi_a - \dot{\psi}_a}{\psi_I} + c_H \psi_a - \zeta c_T, \tag{50}$$

which after dividing through by c_a gives the growth rate of the MAC,

$$\frac{\dot{c_a}}{c_a} = r + \underbrace{\frac{c_H\psi_a + \frac{\psi_H}{\psi_I}\psi_a - \frac{\dot{\psi_a}}{\psi_I}}{c_a}}_{\textit{Endog. Future Gain}} - \underbrace{\frac{\zeta c_T}{c_a}}_{\textit{Damages}}. \tag{51}$$

Peak warming

Lemma 1 *Assume that from the point at which peak warming is reached the marginal abatement cost and damage functions are static. Then TC decreases peak warming.*

Proof: If the MAC and damage functions are static from the point at which peak warming is reached, peak warming is also the steady state. If we conjecture that $\dot{P} = 0$, the MAC, temperature and hence marginal damages will be constant, Equation (49) is satisfied and has the solution

$$MAC^* = \frac{-\zeta c_T^*}{r}. \tag{52}$$

Since a model with TC has a lower MAC in the steady state, from (52) $-c_T^*$ is lower. And since marginal damages are increasing in temperature, the steady state temperature is lower. ■

This lemma depends on static MACs and marginal damages after peak warming, which ensure peak warming is the model's steady state. However, peak warming may not be a steady state. One special case is where both the MAC and damage functions are proportional to consumption. In this case, r in (52) is replaced by $r - g$ and the claim still holds, in fact.

Moreover, in the specialised model of Appendix B, we relax the assumption of no TC after peak warming, allowing TC to continue after peak warming such that it is optimal to cool the earth after peak warming with negative emissions. Yet, we will show that TC still decreases peak warming under highly plausible conditions.

Lemma 1 underpins Remarks 1-3 in the main text as they relate to optimal emissions and temperatures in the long run.

Exogenous TC

Lemma 2 *In a cost-benefit setting, exogenous TC ($\psi_a = 0$) results in a lower initial MAC, lower initial abatement, and a lower initial MAC growth rate than a model without TC.*

Proof: To show this, conjecture that the model with exogenous TC has the same initial MAC as the model without TC. From Equation (51), this implies that initial MAC growth is also the same (the initial conditions ensure marginal damages are the same). The same MAC growth rate combined with a decreasing MAC curve in the model with exogenous TC will lead to faster abatement and lower temperatures after the start. This will lead in turn to faster growth of the MAC in later periods (from Equation 51). From (49), lower emissions and temperatures lead to a lower MAC, which contradicts our conjecture. As a result, the initial MAC needs to be lower. If the MAC is lower and the MAC function is identical at time zero, initial abatement will be lower. Regarding the initial MAC growth rate, (51) shows that since c_T is identical from the initial condition on temperature and c_a is lower, the change in the MAC will initially be lower. ■

In later periods, the lower MAC is compensated by lower marginal damages and the effect of exogenous TC on the growth rate is ambiguous.

Lemma 2 provides the basis for the rest of Remark 1 in the main text, determining the qualitative effect of exogenous TC in abatement technologies on MACs, emissions and temperatures in the short run. In a model of exogenous TC, the MAC and the carbon price are equivalent.

Endogenous, learning-based TC

Corollary 1 *In a cost-benefit setting, endogenous TC ($\psi_a > 0$) has ambiguous effects on the initial MAC, initial abatement and the initial MAC growth rate compared to a model without TC.*

In a model with endogenous TC, there is the additional marginal learning gain (Equation 49), which increases the MAC, all else being equal. Hence, the effect of endogenous TC on the initial MAC, initial abatement and the initial MAC growth rate is in general ambiguous. This provides the remaining basis for Remark 3 in the main text, determining the qualitative effect of endogenous, learning-based TC on optimal MACs, emissions and temperatures in the short run.

A.2 Cost-effectiveness analysis

In the case of cost-effectiveness analysis, damages are replaced by a temperature ceiling $T \leq \bar{T}$, which corresponds to a constraint on cumulative emissions $\bar{T} = \zeta \bar{S}$. The Hamiltonian in Equation (43) is replaced by the following Lagrangian,

$$\mathcal{L} = u\left(c\left(a, H, I, t\right)\right) - \lambda\left(P_{BAU} - a\right) + \mu\psi\left(t, H, I, a\right) - \theta\left(P_{BAU} - a\right), \tag{53}$$

where the Lagrange multiplier θ indicates that whenever $S = \bar{S}$, emissions cannot be positive. The FOCs are

$$u_c c_a + \mu\psi_a = \lambda + \theta, \tag{54}$$

$$\dot{\lambda} = \left(\delta - n\right)\lambda, \tag{55}$$

$$\begin{aligned} &\text{If} \quad S = \bar{S}: \;\; \theta > 0;\; \dot{\theta} \leq 0;\; P \leq 0. \\ &\text{If} \quad S < \bar{S}: \;\; \theta = 0. \end{aligned} \tag{56}$$

Equations (62) and (64) remain the same. Call $\bar{t}$ the time when the constraint hits. Before $\bar{t}$, the Lagrange multiplier is zero and the integral expression for the MAC is

$$\underbrace{c_{a_t}}_{MAC} = \underbrace{\lambda_0 e^{rt}}_{Hotelling} + \underbrace{\overbrace{\psi_{a_t}}^{knowledge\ increment} \overbrace{\int_t^{\infty} e^{-\int_t^u (r+\psi_H)ds} c_{H_u} du}^{...and\ its\ effect\ on\ abatement\ costs}}_{Endog.\ Future\ Gain}. \tag{57}$$

At time $\bar{t}$, the continuity of the costate variable implies continuous MACs and therefore continuous emissions at zero. Using boundary conditions at time $\bar{t}$, $a_{\bar{t}} = P_{BAU}$, $S_{\bar{t}} = \bar{S}$, $\int_0^{\bar{t}} P dt = \bar{S} - S_0$, and assuming that endogenous future learning gains are negligible after time $\bar{t}$, we establish that λ_0 is the present value of the MAC at zero emissions at time $\bar{t}$, $\lambda_0 = e^{-r\bar{t}}\ c_{a_{\bar{t}}}|_{a=P_{BAU}}$ (this requires the knowledge stock at time $\bar{t}$, so there will be no closed-form solution in the general case).

For green R&D or exogenous TC, we have $\psi_a = 0$ and the Hotelling rule is preserved. For general cases of endogenous TC, the Hotelling rule is no longer valid.

Exogenous TC

Lemma 3 *In a cost-effectiveness setting, exogenous TC has no effect on the growth rate of the MAC (the Hotelling rule is unaffected), but results in a lower MAC over the entire path, less initial abatement, higher abatement later, and peak warming is reached earlier.*

Proof: From Equation (51) and the assumptions of cost-effectiveness ($c_T = 0$) and no endogenous TC ($\psi_a = 0$), we see that the MAC increases according to the Hotelling rule at rate r in both models. Conjecture that the initial MAC would start at the same level. The model with TC has lower abatement costs after time zero and would have weakly higher abatement over the whole path, reaching zero emissions earlier. This would violate the condition that cumulative emissions before reaching zero emissions must be equal in both models ($\zeta \int_0^\infty E = \bar{T}$). Hence, the MAC must start lower. Since the MAC function is identical at the start, this must result in lower abatement at the start. Since cumulative emissions must be identical, the abatement path must cross the no-TC abatement path and lead to zero emissions earlier. ■

Endogenous, learning-based TC

Lemma 4 *In a cost-effectiveness setting, endogenous TC ($\psi_a > 0$) has ambiguous effects on the initial MAC, initial abatement and the initial MAC growth rate compared to a model without TC.*

Proof: In a model with endogenous TC, there is the additional endogenous future gains component (Equation 49), which increases the MAC, all else being equal. Hence the effect of endogenous TC on the initial MAC, and in turn initial abatement and the initial MAC growth rate, is in general ambiguous. ■

A.3 Technological processes with different learning speeds

It is straightforward to extend the model to more technologies. We will show the derivation for two families of abatement technology, each with a different TC process. Assume the same conditions as above apply to each technology, $c_a < 0; c_{aa} < 0; c_a|_{a<0} = 0$, meaning that within each abatement family there are diminishing returns to scale (see Bramoullé and Olson, 2005, for a model with constant MACs and a technology accumulation function $\dot{H} = a$). The assumption of differentiability and zero MACs for negative emissions implies zero MAC at zero abatement. This will ensure an interior solution for each technology family: at least some abatement is optimal for each technological family.[24]

[24] In case abatement costs are non-continuous at zero, a non-negativity condition $a_i \geq 0$ should be added and will lead to the obvious conclusion that families of technologies that are more expensive than alternative families

The planner's objective is

$$\max_{a,I} \int_0^\infty e^{-(\delta-n_t)t} u\left(c\left(a_1, a_2, H_1, H_2, I_1, I_2, T(S), t\right)\right) dt, \tag{58}$$

subject to

$$\dot{S} = P_{BAU} - a_1 - a_2; \dot{H}_i = \psi\left(t, H_i, I_i, a_i\right). \tag{59}$$

The current value Hamiltonian of the problem is

$$\mathcal{H} = u\left(c\left(a_1, a_2, H_1, H_2, I_1, I_2, T(S), t\right)\right) - \lambda\left(P_{BAU} - a_i\right) + \sum_{i=1,2} \mu_i \psi^i\left(t, H_i, I_i, a_i\right). \tag{60}$$

The seven FOCs for $i = 1, 2$ include

$$u_c c_{a_i} + \mu_i \psi^i_{a_i} = \lambda, \tag{61}$$

$$u_c = \mu_i \psi^i_{I_i}, \tag{62}$$

$$\dot{\lambda} = (\delta - n_t)\lambda - u_c \zeta u_T, \tag{63}$$

$$\dot{\mu_i} = (\delta - n_t)\mu_i - u_c c_{H_i} - \mu \psi_{H_i}. \tag{64}$$

Integrating the costate equations (64) for each technology family and substituting in (61) allows us to write (12) for each technology family separately.

Appendix B Specialised model

To obtain further analytical results, we must specialise some functional relationships. The assumptions regarding welfare and warming are the same as in the preceding section. We now assume emissions are produced according to Equation (4) in the main text, and we assume quadratic damages and total abatement costs, like Equations (3) and (5) in the main text, respectively. Exogenous TC that reduces abatement costs over time is represented by the slope of the MAC curve φ decreasing as a function of time.[25] Instead of a general knowledge stock, we assume that endogenous, learning-based TC is specifically driven by cumulative abatement $A_t = \int_{-\infty}^t a_u du$.[26]

We can write cumulative emissions as a function of time and the state variable S:

$$H(t, S) = A = \int_0^t P_{BAU_u} du - S + S_0. \tag{65}$$

of technologies, even at very low levels of deployment, and even when endogenous future gains are taken into account, should not be deployed yet.

[25] The slope of the MAC curve is also affected by a separate R&D knowledge stock, which is accumulated by R&D investments (but not by abatement), $\dot{\tilde{H}} = \tilde{\psi}\left(\tilde{I}, t\right)$. This stock is optimized according to a third FOC, $u_c = \tilde{\mu}\tilde{\psi}_{\tilde{I}}$.

[26] This is also compatible with a knowledge stock which builds up proportional to abatement, and a scaling factor which depends on both time and deployment-dependent R&D investments $\dot{H} = \varrho\left(I, t\right) a$. However, in this case we assume that optimal investments make the function $\varrho\left(I^*, t\right)$ constant over time (the full model would have a second optimality condition $u_c = \mu \rho_I a$). We normalise the unit of H such that $\varrho = 1$.

This model can be solved with only one state variable and its shadow price λ incorporates both damages and endogenous, learning-based TC.

We obtain the following expression for consumption per capita:

$$c = c_0 \exp\left[gt - \frac{\varphi_t}{2\sigma_t^2 c^2} a^2 (A/A_0)^{-\chi} - \frac{\gamma}{2}\zeta^2 S^2\right]. \tag{66}$$

Note that we have written the emissions control rate as a function of abatement $\mu = \frac{a}{\sigma_t c}$. The MAC is

$$c_a = -\frac{\varphi_t}{\sigma_t^2 c} a (A/A_0)^{-\chi} = -c\frac{\varphi_t}{P_{BAU_t}^2} a (A/A_0)^{-\chi}. \tag{67}$$

Utility $u(c) = \frac{c^{1-\eta}}{1-\eta}$. We assume decreasing population growth $n = n_0 e^{-g_n t}$ and standardise initial population to one. The welfare functional is

$$\max \int_0^\infty e^{-(\delta - n_t)t} u \; dt. \tag{68}$$

The present Value Hamiltonian is

$$H^{PV} = e^{-(\delta-n_t)t} u\left(c\left(t, S, a\right)\right) - \lambda^S \left(P_{BAU} - a\right). \tag{69}$$

The FOCs are

$$\lambda^S = -e^{-(\delta-n_t)t} u_c c_a, \tag{70}$$

$$\dot{\lambda^S} = e^{-(\delta-n_t)t} u_c c_S, \tag{71}$$

where $c_S = c\left(-\gamma\zeta^2 S - \frac{\chi\varphi_t}{2A_0\sigma_t^2 c^2} a^2 (A/A_0)^{-\chi-1}\right)$ includes both marginal damages and endogenous learning gains from TC.

Integrate (71) between time t and infinity with terminal condition $\lim_{t\to\infty}\lambda^S = 0$:

$$\lambda_t^S = \int_t^\infty e^{-(\delta-n_u)u} c_u^{1-\eta}\left(\gamma\zeta^2 S + \frac{\chi\varphi_t}{2A_0}\mu^2 (A/A_0)^{-\chi-1}\right) du. \tag{72}$$

Combining this result with (70) and dividing by $-e^{-(\delta-n_t)t} u_{c_t}$ gives the expression MAC=SCC:

$$-c_a = \int_t^\infty e^{-\delta(u-t)+(n_u u - n_t t)} \left(\frac{c_u}{c_t}\right)^{-\eta} c_u \left(\gamma\zeta^2 S + \frac{\chi\varphi_u}{2A_0}\mu^2 (A/A_0)^{-\chi-1}\right) du. \tag{73}$$

Acknowledging $\ln\frac{c_u}{c_t} = \int_t^u d\ln c_s \Leftrightarrow \frac{c_u}{c_t} = e^{\int_t^u \frac{\dot{c}}{c}ds} \Leftrightarrow \left(\frac{c_u}{c_t}\right)^{-\eta} = e^{-\eta\int_t^u \frac{\dot{c}}{c}ds}$ shows that the first factor in the integral is the discount factor with the Ramsey discount rate $\delta - n_t + \eta\frac{\dot{c}}{c}$.

Note that when we write the MAC as a function of BAU emissions, both MACs and damages are proportional to consumption, so we can write this equation relative to consumption while reducing the discount rate by the growth rate of consumption:

$$\frac{\varphi_t}{P_{BAU_t}^2} a (A/A_0)^{-\chi} = \int_t^\infty e^{-\delta(u-t)+(n_u u - n_t t)-(\eta-1)\int_t^u \frac{\dot{c}}{c}ds}\left(\gamma\zeta^2 S_u + \frac{\chi\varphi_u}{2A_u}\mu_u^2 (A_u/A_0)^{-\chi}\right) du. \tag{74}$$

To find a differential equation for the MAC, first take the time derivative of (70):

$$\dot{\lambda^S} = -e^{-(\delta-n_t)t}(-\delta + n_t + \dot{n}_t t)u_c c_a - e^{-(\delta-n_t)t}\dot{u}_c c_a - e^{-(\delta-n_t)t}u_c\dot{c}_a. \tag{75}$$

Substitute out $\dot{\lambda^S}$ from (71) and (75), divide by $e^{-(\delta-n_t)t}u_c$, and use $-\frac{\dot{u}_c}{u_c} = \eta\frac{\dot{c}}{c}$ to give

$$-\dot{c_a} = \left(\delta - n_t - \dot{n}_t t + \eta\frac{\dot{c}}{c}\right)(-c_a) + c_S, \tag{76}$$

which corresponds to

$$\underbrace{\frac{\dot{c_a}}{c_a}}_{MAC\ growth} = r - \underbrace{\frac{\zeta c_T}{c_a}}_{Marginal\ Damages} + \underbrace{\frac{c_A}{c_a}}_{Endogenous\ Future\ Gains}. \tag{77}$$

To find a differential equation for abatement, we take the time derivative of (67) to calculate the growth rate of the MAC :

$$\frac{\dot{c_a}}{c_a} = \frac{\dot{\varphi}}{\varphi} + \frac{\dot{a}}{a} - 2\frac{\dot{\sigma}}{\sigma} - \frac{\dot{c}}{c} - \chi\frac{a}{A}.$$

Plugging the growth rate of the MAC into (77) results in (78) (extended to decreasing population growth), i.e.,

$$\frac{\dot{a}}{a} = \underbrace{\delta - n_t - \dot{n}_t t + (\eta+1)\frac{\dot{c}}{c}}_{r+g} - \frac{\dot{\varphi}}{\varphi} + 2\frac{\dot{\sigma}}{\sigma} + \frac{\chi a}{A} - \frac{1}{2}\frac{\chi a}{A} - \frac{\sigma^2 c^2 \gamma \zeta^2 S}{\varphi a\,(A/A_0)^{-\chi}}. \tag{78}$$

Note that if BAU emissions are constant, we have $\frac{\dot{\sigma}}{\sigma} = -g$, which boils down to using $r - g$ instead of $r + g$.

Equation (78) shows that TC will increase initial abatement speed for real-world parameters. Let us first show this for exogenous TC. Denote the (negative) growth rate of exogenous TC as $g_{TC} = \dot{\varphi}/\varphi$ and define $\frac{\partial a}{\partial g_{TC}}$ as the difference in abatement between two identical models with marginally different TC growth rates. Similarly, define $g_a = \dot{a}/a$ and denote $\frac{\partial a}{\partial g_{TC}}$ as the difference in abatement between two models, where the TC growth rate has been marginally altered. A change in the exogenous TC rate will affect initial abatement, but not A, S and φ, because they are defined by their initial conditions. Taking the derivative of Equation (78) with respect to g_{TC} at time zero gives

$$-\frac{\partial g_a}{\partial g_{TC}} = 1 - \frac{\gamma\zeta^2 S}{\varphi a(A/A_0)^{-\chi}}\frac{\partial a/a}{\partial g_{TC}}. \tag{79}$$

The factor $\frac{\gamma\zeta^2 S}{\varphi a(A/A_0)^{-\chi}}$ is the marginal damage over the MAC, which is typically lower than 2%, (it is the adjusted discount rate $r + g - 2\frac{\dot{\sigma}}{\sigma}$ at the steady state and considerably smaller at the start of the transition). So even if $\frac{\partial a/a}{\partial g_{TC}}$ is large, say initial abatement increases by 5% for a 1% increase of TC, the RHS remains large positive. So exogenous TC increases the initial abatement speed.

For endogenous, learning-based TC, a very similar argument can be made. Denote the (negative) growth rate of endogenous, learning-based TC as $g_{TC} = \frac{\frac{d}{dt}(A/A_0)^{-\chi}}{(A/A_0)^{-\chi}} = -\frac{\chi a}{A}$. Taking the derivative of equation (78) with respect to g_{TC} at time zero gives

$$-\frac{\partial g_a}{\partial g_{TC}} = 1/2 - \frac{\gamma\zeta^2 S}{(\varphi a(A/A_0)^{-\chi}} \frac{\partial a/a}{\partial g_{TC}}. \tag{80}$$

The same term as for exogenous TC reappears and for the same reasons this will be small. Moreover, $\frac{\partial a/a}{\partial g_{TC}}$ is not only small, it is often negative (initial abatement increases with endogenous, learning-based TC), as it is for our parameters. The term 1/2 comes from the fact that the endogenous future gain incentive creates a flattening effect on the abatement path that is half the size of the opposite cost-reduction effect. So again, endogenous, learning-based TC increases the initial abatement speed.

The extension to a market economy is a straightforward application of the first welfare theorem. The first order conditions will be the same provided that λ^A equals the sum of the tax and the subsidy. Appendix D shows a comparison between the market equilibrium and the social optimum in a more elaborate model.

Peak warming

We have already developed several results for the general model in Appendix A. These are of course also valid in this more specialised model. However, by adding more structure, we can produce a few more results.

Lemma 5 *TC decreases peak warming if at the time of peak warming optimal abatement satisfies*

$$\frac{\left(r + g - 2\frac{\dot\sigma}{\sigma} \overbrace{-\frac{\dot a}{a}}^{neg} \overbrace{-\frac{\dot\varphi}{\varphi} + \frac{\chi a}{2A}}^{pos}\right)}{\left(r + g - 2\frac{\dot\sigma}{\sigma}\right)} \frac{\varphi_{t^*}\left(\frac{A}{A_0}\right)^{-\chi}}{\varphi_0} < 1. \tag{81}$$

Proof: The claim follows from rewriting Equation (78) at the time of peak warming as

$$\gamma\zeta T^* = \left(r + g - 2\frac{\dot\sigma}{\sigma} \overbrace{-\frac{\dot a}{a} - \frac{\dot\varphi}{\varphi} + \frac{\chi a}{2A}}^{-\frac{d}{dt}\left(\frac{MAC}{c}\right) - \frac{c_A}{c}} \right) \varphi a \left(\frac{A}{A_0}\right)^{-\chi}, \tag{82}$$

and dividing this by the equivalent expression for the model without TC. ■

This condition is satisfied unless peak warming is reached implausibly quickly. The first factor on the left-hand side of the inequality converges to one in the long term and tends to be very close to one at peak warming. The numerator comprises the discount rate, plus the degrowth rate of the MAC (the cost-reduction effect, $-d/dt(MAC/c)$), minus the endogenous future gains of learning-based TC. In the static model, this is just $r+g-2\frac{\dot\sigma}{\sigma}$, but in the presence of TC we have additional terms. The second factor is much smaller than one. It is the relative reduction of the slope of the MAC curve due to TC at the time of peak warming.

Lemma 5 extends Lemma 1 to the case where TC continues after peak warming is reached. Therefore, it provides a more general set of circumstances underpinning Remarks 1-3 as they relate to optimal policy in the long run.

Endogenous, learning-based TC under cost-effectiveness analysis

Lemma 6 *In a cost-effectiveness setting, compared to a model without TC, the model with endogenous TC will have lower MAC growth, a lower MAC over the entire path, less initial abatement, higher later abatement, and reach peak warming earlier.*

Proof: The emissions path will still be steeper, because the cost-reduction effect $\frac{\chi a}{A}$ dominates the endogenous future gain effect $\frac{\chi a}{A}$ in Equation 78 (where $\gamma = 0$). Since cumulative emissions are identical with and without TC in the cost-effectiveness case, the emissions paths must cross each other, with emissions higher at the start under endogenous TC, lower at the end, and zero emissions reached earlier. Higher initial emissions implies a lower initial MAC. Equation (77) shows that endogenous future gains will reduce the growth rate of the optimal MAC (marginal damages are zero in the cost-effectiveness analysis). A lower initial MAC combined with lower growth implies a lower MAC over the entire path. ■

Note that the shadow price of carbon, λ in Equation (55), always follows a Hotelling path irrespective of TC. This result is also reported in Goulder and Mathai (2000).

Decreasing carbon intensity (AEEI)

Lemma 7 *In a cost-effectiveness setting, compared to a model with constant carbon intensity ($\frac{\dot{\sigma}}{\sigma} = 0$), the model with decreasing carbon intensity will have lower initial emissions and reach peak warming later, if optimal initial emissions are less than half of BAU emissions.*

Proof: To investigate the role of decreasing carbon intensity $\dot{\sigma}$ on the emission path, we need to find an expression for the change in emissions, because abatement depends on σ. Taking the time derivative of abatement $a = \sigma c - P$ yields

$$\frac{\dot{a}}{a} = -\frac{1-\mu}{\mu}\frac{\dot{P}}{P} + \frac{1}{\mu}\underbrace{\left(\frac{\dot{\sigma}}{\sigma} + \frac{\dot{c}}{c}\right)}_{\frac{\dot{P}_{BAU}}{P_{BAU}}}. \tag{83}$$

Combining this result with (78) and acknowledging that the damage term is zero for cost-effectiveness analysis (replaced by a constraint on temperature), gives

$$\frac{\dot{P}}{P} = g + \left(2 - \frac{P_{BAU}}{P}\right)\frac{\dot{\sigma}}{\sigma} + \left(\frac{P_{BAU}}{P} - 1\right)\left(\frac{\dot{\varphi}}{\varphi} - \frac{\chi a}{2A} - r\right), \tag{84}$$

where $P_{BAU} = \sigma c$. [27]

[27] When we assume constant exogenous TC ($\frac{\dot{\varphi}}{\varphi} = g_{\varphi}$.) and decreasing carbon intensity at a constant rate ($\frac{\dot{\sigma}}{\sigma} = g_{\sigma}$), we can integrate $\dot{P}$ using boundary conditions $\int_0^T P_t dt = \bar{S}$ and $P_T = 0$. This gives

$$P_t = P_{BAU_0} e^{(g+g_{\sigma})t} - a_0 e^{(g+2g_{\sigma}-g_{\varphi}+r)t}. \tag{85}$$

The factor $2-\frac{P_{BAU}}{P}$ is positive whenever emissions are less than half of BAU emissions, $(P<\frac{1}{2}P_{BAU})$. So if initial emissions are less than half of BAU emissions, the entire emissions path will be flatter. Since the exogenous temperature constraint sets cumulative emissions , the flatter emissions path will lead to lower initial emissions and later peak warming. ■

By contrast, for lenient, unambitious climate scenarios, AEEI will initially steepen the emissions path (when $P_t>\frac{1}{2}P_{BAU_t}$) and later on flatten the emissions path (when $P_t<\frac{1}{2}P_{BAU_t}$). The effect of AEEI on initial emissions and the time of peak warming is therefore ambiguous.

The case of cost-benefit analysis is more complicated because AEEI increases peak warming. Consider again the case where $P_0=\frac{1}{2}P_{BAU_0}$. The growth rate of emissions at time zero is

$$\frac{\dot{P}}{P}=g+\frac{\dot{\varphi}}{\varphi}-r+\overbrace{c\gamma\zeta^2 S_0\frac{\sigma}{\varphi\mu}}^{c_S/c_E}.$$

This is unaffected by AEEI because marginal damages are set by initial cumulative emissions and the MAC set by our hypothesis that $\mu=1/2$. In subsequent periods, AEEI will again flatten the emissions path. This flattening will be reinforced by an increasing MAC in later periods (effect of σ in the last term). However, combined with larger cumulative emissions, initial emissions are ambiguous. Nevertheless, in the limit case where initial emissions are extremely low (approach zero), AEEI will flatten the emissions path and therefore (slightly) decrease initial emissions.[28] By contrast, in the limit case where initial emissions approach BAU emissions, AEEI reduces the emission speed $\lim_{P->P_{BAU}}\dot{P}=g+\frac{\dot{\sigma}}{\sigma}$.

These results are summarised in Remark 2 in the main text.

Heterogenous learning rates

A more detailed model may have several groups of abatement technologies, each with a MAC function $c_{i,t}\frac{\varphi_{i,t}}{P^2_{BAU_t}}a_i\left(\frac{A_i}{A_{0i}}\right)^{\chi_i}$. Cumulative emissions are now $S_t=S_0+P_{BAU}t-\sum A_i$. For N groups of technologies, the model now has N decision variables and N stock variables (cumulative abatement for each group of technologies). Assuming a constant discount rate, the integral form of the Euler equations for each technology is

$$\underbrace{\frac{\varphi_{i,t}}{P^2_{BAU_{i,t}}}a_{i,t}\left(A_{i,t}/A_{0,i}\right)^{-\chi_i}}_{MAC_i^{\%}}=\int_t^{\infty}\underbrace{e^{-(r-g)t}}_{Discount\ factor}\left(\underbrace{\gamma\zeta^2 S_u}_{Marg\ damages^{\%}}+\underbrace{\frac{\chi_i\varphi_{i,u}}{2A_{i,u}}\mu^2_{i,u}\left(A_{i,u}/A_{i,0}\right)^{-\chi_i}}_{Endogenous\ future\ gains_i^{\%}}\right)du. \tag{86}$$

The first term describes BAU emissions over time and the second term describes abatement over time. AEEI has a double effect: it reduces BAU emissions (steepening the emissions path) and it slows down abatement (flattening the emissions path).

[28] At low emissions, the equation of motion for emissions converges to $\lim_{E->0}\dot{E}=c_0\sigma_0\left(\frac{\dot{\varphi}}{\varphi}-\frac{\dot{\sigma}}{\sigma}-r+\frac{\sigma_0}{\varphi_0}\gamma\zeta^2 S_0\right)$ showing that fast AEEI leads to a flatter emissions path.

Appendix C The specialised model with inertia

This model builds on the specialised model described in Appendix B. It is the model that we will use for quantitative analysis. To fit the model to the observation that BAU emissions in the IPCC-NGFS database are constant (or at least not increasing), we set AEEI equal to GDP growth. This also simplifies notation.

Call $v = \dot{a}$ the abatement speed. Assume a quadratic penalty on abatement speed (as a proportion of GDP) of $\frac{\theta}{2}v^2$. Consumption per capita now becomes

$$c = c_0 e^{\left(gt - \frac{\varphi_t}{2}a^2(A/A_0)^{-\chi} - \frac{\theta}{2}v^2 - \frac{\gamma}{2}\zeta^2 S^2\right)}. \tag{87}$$

The present value Hamiltonian is

$$H^{PV} = e^{(-\delta+n_t)t}u(c) - \lambda^S(E_{BAU} - a) + \lambda^a v, \tag{88}$$

with FOCs

$$\lambda^a = e^{(-\delta+n_t)t}u_c c\theta v, \tag{89}$$

$$\dot{\lambda^a} = e^{(-\delta+n_t)t}u_c c\varphi_t a\,(A/A_0)^{-\chi} - \lambda^S, \tag{90}$$

$$\dot{\lambda^S} = e^{(-\delta+n_t)t}u_c c_S. \tag{91}$$

Differentiate the FOC of the maximisation:

$$\dot{\lambda^a} = e^{(-\delta+n_t)t}u_c c\theta v\left[-\delta + n_t + \dot{n}t - \eta\frac{\dot{c}}{c} + \frac{\dot{c}}{c} + \frac{\dot{v}}{v}\right], \tag{92}$$

substitute this result in Equation (90) and divide by $e^{(-\delta+n_t)t}u_c c\theta$:

$$v\left[-\delta + n_t + \dot{n}_t t - (\eta-1)\frac{\dot{c}}{c}\right] + \dot{v} = \frac{\varphi_t a\,(A/A_0)^{-\chi}}{\theta} - \frac{\lambda^S}{e^{(-\delta+n_t)t}u_c c\theta}, \tag{93}$$

with

$$n_t + \dot{n}t = n_0 e^{-g_n t}\,(1 - g_n t)\,, \tag{94}$$

and

$$\frac{\dot{c}}{c} = g + \frac{(\varphi_0 - \varphi_\infty)\,g_\varphi}{2}a^2\,(A/A_0)^{-\chi} - \varphi_t a v\,(A/A_0)^{-\chi} + \frac{\varphi_t}{2}a^2\chi\,(A/A_0)^{-\chi}\,\frac{a}{A} - \theta v\dot{v} - \gamma\zeta^2 S\,(E_{BAU} - a)\,. \tag{95}$$

The growth rate of consumption is very close to g, but the component $-\theta v\dot{v}$ cannot be neglected. Reorganise to obtain a differential equation in $\dot{v}$,

$$\dot{v} = \frac{1}{1+(\eta-1)\theta v^2}\left[\left(\delta - n_0 e^{-g_n t}\,(1-g_n t) + (\eta-1)\overbrace{\left(\frac{\dot{c}}{c}\right)}^{Net\ of\ \theta v\dot{v}}\right)v + \frac{\varphi_t a\,(A/A_0)^{-\chi}}{\theta} - \frac{\lambda^S}{e^{(-\delta+n_t)t}u_c c\theta}\right]. \tag{96}$$

We now have a system of four differential equations in four variables S, a, v, λ^S (Equations 40, 78, 96 and 91). The boundary conditions are

$$\begin{aligned} S(0) &= S_0, \\ a(0) &= a_0 = E_{BAU} - E_0, \\ a(\infty) &= E_{BAU}, \\ v(\infty) &= 0. \end{aligned} \tag{97}$$

The MAC is the current value shadow price of carbon, expressed in consumption units, i.e.,

$$\text{MAC} = \frac{\lambda^S e^{\delta t}}{u_c}. \tag{98}$$

From Equation (93) we can rewrite this as [29]

$$\text{MAC} = e^{n_t t} \left\{ \underbrace{c\varphi_t a \left(A/A_0\right)^{-\chi}}_{\partial c/\partial a \text{ standard MAC}} \underbrace{+ c\theta v \left[\delta - n_t - \dot{n}_t t + (\eta - 1)\frac{\dot{c}}{c}\right] - c\theta \dot{v}}_{\text{Abatement speed costs (pos)}} \right\}. \tag{99}$$

Alternatively, the integration of Equation (91) gives the MAC as the sum of both the SCC and the endogenous future gains:

$$\text{MAC} = \frac{\lambda^S e^{\delta t}}{u_c} = \int_t^\infty e^{n_u u} \underbrace{e^{-\delta(u-t) - \eta \int_t^u \frac{\dot{c}}{c} ds}}_{\text{Discount factor}} c_u \left(\gamma\zeta^2 S + \frac{\chi\varphi_u}{2A} a^2 \left(A/A_0\right)^{-\chi}\right) du. \tag{100}$$

The discount factor is the standard Ramsey discount factor as can be seen from $\left(\frac{c_u}{c_t}\right)^{-\eta} = e^{-\eta \int_t^u \frac{\dot{c}}{c} ds}$.

Appendix D Model with energy inputs and knowledge spillovers from private to aggregate

In this Appendix, we extend the simple model of Section 2 in two directions. First, we explicitly model energy inputs, giving the model two mechanisms for abatement: energy efficiency and carbon intensity of energy, which is modelled with a simple abatement variable rather than explicit clean and dirty energy. We assume that innovation affects the carbon intensity of energy and not energy efficiency. This feature creates a different role for the carbon price and the innovation subsidy: the optimal tax applies to all abatement technologies and equalises the MAC between sectors, whereas subsidies target the sectors with the largest innovation potential.

Second, we assume that when firms innovate, they obtain an individual advantage before their knowledges diffuses in the wider economy. That is, not all knowledge is shared with other producers, some knowledge is private as in Gerlagh et al. (2009); Greaker and Pade (2009). We

[29] For analytical simplicity, we usually focus on the marginal effect of abatement on consumption per capita, i.e. c_a. However, since a is expressed as worldwide abatement, the marginal effect of abatement on total consumption is more relevant, because it corresponds to costs in production. So we report the latter and we multiply costs by population size $e^{n_t t}$.

show that the optimal R&D subsidy does not compensate for private knowledge, only for the gradual leakage from private firms to public firms.

Firm i produces with the following Cobb-Douglas production function, combining labour L, capital K and energy E (we leave out index i for now),

$$Y = (A_L L)^{1-\nu-\alpha} K^{\alpha} E^{\nu} \Gamma \Lambda, \tag{101}$$

where Γ and Λ represent abatement costs and damages from climate change, respectively. Emissions are defined as $P = \psi(1-\mu)E$, where $\mu \in [0,1]$ is abatement effort and ψ is the carbon intensity of energy inputs. Dirty inputs are expressed in tonnes of CO_2. Companies can abate $a = \mu\psi E$, which is costly. Abatement costs are proportional to production and quadratic in abatement,

$$\Gamma = e^{-\frac{\varphi}{2}a^2\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi}}, \tag{102}$$

where H is the public knowledge stock and h is the private knowledge stock. Companies make an effort z to reduce their abatement costs, building up their private knowledge stock. However, the private knowledge stock becomes public after some time. This leads to the following dynamics: $\dot{h}_i = z_i - \delta h_i$ and $\dot{H} = \sum_i^N \delta h_i$. The knowledge stock reduces abatement costs with a constant elasticity χ. Effort z costs $C(z)$.

As TC is the result of R&D which is different from abatement, the optimal policy requires two instruments, a Pigouvian carbon tax τ and a knowledge subsidy s which is proportional to z. Also, the model has two different abatement mechanisms: energy efficiency and carbon intensity of energy. TC only applies to the latter. As a result, using only a carbon tax leads to excess energy efficiency effort and insufficient TC, whereas using only a subsidy leads to insufficient energy efficiency.

Market equilibrium

We assume the total number of companies N is large enough that companies are price takers and the market is competitive. Normalising the price of the final good to one, the profit function of the company is

$$\pi = A_L L^{1-\nu-\alpha} K^{\alpha} E^{\nu} \Gamma \Lambda - wL - rK - p_E E - C(z) - \tau\psi\left(1-\mu\right)E + sz. \tag{103}$$

s.t. $\dot{H} = \sum_i^N \delta h_i; \dot{h}_i = z_i - \delta h_i$. To maximize profits inter temporally, the company solves the following current value Hamiltonian

$$\mathscr{H}^{CV} = \pi + \lambda^H \left(\sum_i^N \delta h_i\right) + \lambda^{h_i}\left(z_i - \delta h_i\right). \tag{104}$$

The FOCs include

$$\frac{\partial \mathscr{H}}{\partial \mu} = 0 \Leftrightarrow Y\varphi a \left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi} = \tau, \tag{105}$$

$$\frac{\partial \mathscr{H}}{\partial E} = 0 \Leftrightarrow \nu Y/E = \psi\left(1-\mu\right)\tau, \tag{106}$$

$$\frac{\partial \mathscr{H}}{\partial z} = 0 \Leftrightarrow \frac{dC}{dz} - s = \lambda^h, \tag{107}$$

$$\dot{\lambda^H} = r\lambda^H - Y\frac{\varphi}{2}a^2\chi\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi-1}\frac{1}{H_0+\omega h_0}, \tag{108}$$

$$\dot{\lambda^h} = (r+\delta)\,\lambda^h - Y\frac{\varphi}{2}a^2\chi\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi-1}\frac{\omega}{H_0+\omega h_0} - \delta\lambda^H. \tag{109}$$

The integral form of (108) and (109) gives

$$\lambda_t^H = \int_t^{\infty} e^{-r(u-t)}\left[Y_i\frac{\varphi}{2}a^2\chi\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi-1}\frac{\omega}{H_0+\omega h_0}\right]du, \tag{110}$$

$$\lambda_t^h = \int_t^{\infty} e^{-(r+\delta)(u-t)}\left[Y_i\frac{\varphi}{2}a^2\chi\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi-1}\frac{\omega}{H_0+\omega h_0} + \delta\lambda^H\right]du. \tag{111}$$

The company sets the subsidy net of R&D costs equal to λ^h, which is the value of knowledge. This value of knowledge is composed of all future direct effects on its abatement efforts (the first term in the integral) plus the fact that the firm contributes to aggregate knowledge, which will also make abatement cheaper for the individual firm. Yet, the shadow price of this aggregate knowledge λ_t^H corresponds to the aggregate productivity gains for firm i only, which is N times lower than the gain from a social point of view. To make that point more formally, we now derive the social optimum.

Social optimum

We assume taxes and subsidies are collected lump sum, without a cost of public funds. All variables are now aggregate quantities, unless indexed by i. The social planner maximises the present value of production

$$Y = \sum_i^N Y_i = \sum_i^N A_L L_i^{1-\nu-\alpha} K_i^{\alpha} E_i^{\nu} e^{-\frac{\varphi}{2}a_i^2\left(\frac{H+\omega h_i}{H_0+\omega h_0}\right)^{-\chi} - \frac{\gamma}{2}\zeta^2 S^2} - \sum_i^N C(z_i), \tag{112}$$

s.t. $\dot{H} = \sum_i^N \delta h_i; \dot{h_i} = z_i - \delta h_i; \dot{S} = \psi\,(1-\mu)\,E$. The current value Hamiltonian is

$$\mathscr{H}^{CV} = Y + \lambda^H\left(\sum_i^N \delta h_i\right) + \sum_i^N \lambda_i^h\,(z_i - \delta h_i) - \lambda^S\psi\,(1-\mu)\,E. \tag{113}$$

The FOCs include

$$Y\varphi a\left(\frac{H+\omega h_i}{H_0+\omega h_{i_0}}\right)^{-\chi} = \lambda^S, \tag{114}$$

$$\nu Y/E = \psi\,(1-\mu)\,\lambda^S \tag{115}$$

$$\frac{dC}{dz_i} = \lambda_i^h, \tag{116}$$

$$\dot{\lambda^H} = r\lambda^H - \sum_i^N Y_i\frac{\varphi}{2}a^2\chi\left(\frac{H+\omega h_i}{H_0+\omega h_{i_0}}\right)^{-\chi-1}\frac{1}{H_0+\omega h_{i_0}}, \tag{117}$$

$$\dot{\lambda}_i^h = (r+\delta)\lambda_i^h - Y_i \frac{\varphi}{2} a^2 \chi \left(\frac{H+\omega h_i}{H_0+h_{i_0}}\right)^{-\chi-1} \frac{\omega}{H_0+\omega h_{i_0}} - \delta\lambda^H. \tag{118}$$

The integral form of (117) and (118) gives

$$\lambda_t^H = \int_t^\infty e^{-r(u-t)} \left[\sum_i^N Y_i \frac{\varphi}{2} a^2 \chi \left(\frac{H+\omega h_i}{H_0+\omega h_{i_0}}\right)^{-\chi-1} \frac{\omega}{H_0+h_{i_0}}\right] du, \tag{119}$$

$$\lambda_t^h = \int_t^\infty e^{-(r+\delta)(u-t)} \left[Y_i \frac{\varphi}{2} a^2 \chi \left(\frac{H+\omega h_i}{H_0+\omega h_{i_0}}\right)^{-\chi-1} \frac{\omega}{H_0+h_{i_0}} + \delta\lambda^H\right] du. \tag{120}$$

Equation (117) shows that the shadow price of aggregate knowledge is now N times larger. Therefore, the term $\delta\lambda^H$ inside the integral of (120) of the company specific shadow price is N times larger. Indeed the individual company may take into account that spillovers to the aggregate knowledge stock will be beneficial for itself too, but this advantage is N times larger from a social perspective. Comparing both solutions indicates that the subsidy should be

$$s = (1-1/N) \int_t^\infty e^{-(r+\delta)(u-t)} \left[\delta\lambda^H\right] d\tau. \tag{121}$$

Note that when spillovers are negligible the subsidy converges to zero, $\lim_{\delta->0} s = 0$, whereas when spillovers are extremely large, the subsidy converges to the shadow price of the aggregate knowledge stock $\lim_{\delta->\infty} s = (1-1/N)\lambda^H$.

The first FOCs (114) and (115) indicate that both the MAC via decarbonisation of energy and the MAC via energy efficiency gains equal the SCC. Comparing this with the FOCs of the market equilibrium indicates that the optimal tax should be set at the SCC, $\tau = \lambda^S$.

Appendix E Correspondence between exogenous and endogenous TC

E.1 Method of isolating endogenous future gain effect

Our method of isolating the endogenous future gain effect proceeds in two steps. The first step is to estimate a time-dependent MAC function (i.e., exogenous TC), with a MAC that is identical to the endogenous TC model at each point in time. This can be achieved using a sufficiently high-order polynomial. The second step is to take the exogenous TC MAC curve and use it to recalculate optimal model trajectories. Any difference between the optimal paths of the endogenous TC model and its exogenous replica must then be down to the endogenous future gain effect.

We can show this in theory. Call a^* the optimal abatement path of the model with endogenous TC. Conjecture that this is also the optimal path of the model with exogenous TC. By the assumption of identical MACs, the abatement path a^* will result in the same left-hand side of Equation (12) at each point in time. Given that we use the same climate model (Equation 40), the model without TC results in the same temperature path. Figure A1 shows that since MACs are identical, total abatement costs are also identical. Hence, consumption and the discount rate

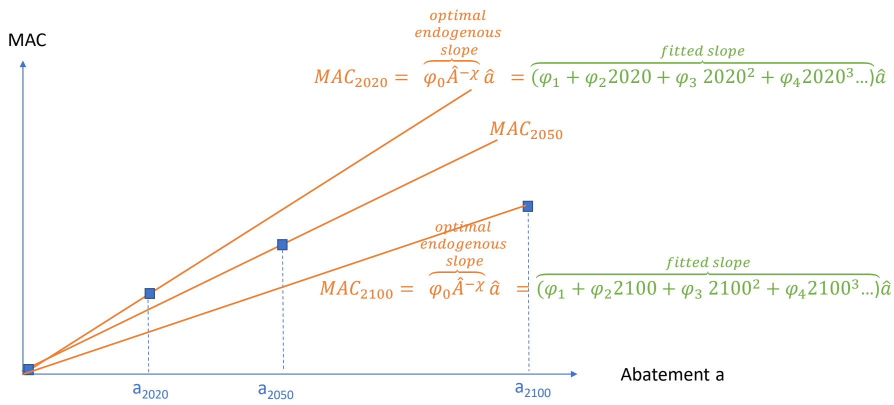

Figure A1: Graphical representation of the comparison between endogenous TC and exogenous TC with identical MAC function. The endogenous TC model has a linear marginal abatement cost function $MAC^{\%} = \varphi A^{-\chi} a$. The optimal path (which depends on cumulative abatement) results in optimal slopes at each point in time. Next, a 25th-degree polynomial in time is fitted to obtain the same MAC but without the dependence on cumulative abatement. $MAC = (\varphi_1 + \varphi_2 t + \varphi_3 t^2 + \varphi_4 t^3 ...)\, a$.

are the same. However, in the exogenous TC model, the absence of the endogenous future gain effect decreases the right-hand side of Equation (12). Hence, our conjectured abatement path a^* does not solve Equation (12), and the initial MAC and abatement are lower in the model with exogenous TC. Lower abatement results in higher temperatures, higher marginal damages, and a higher right-hand side of Equation (12). Therefore, both models converge in the long run. If TC is zero after peak warming ($\forall t > t^{peak} : \dot{H} = 0$), Equation (12) shows that the optimal abatement path of the endogenous TC model also solves the exogenous TC model, i.e., peak warming and long-run temperatures are identical (although peak warming comes earlier in the endogenous TC model).

E.2 Static approximation of MAC function under exogenous TC

We can also test the ability of a suitably calibrated static model to approximate the optimal solution of a model of exogenous TC. Figure A2 shows our approach visually. Exogenous TC results in an optimal pair $\{MAC^{\%*}, a^*\}$ at each moment in time, where again $MAC^{\%} = \varphi_t a$, i.e., the MAC as a proportion of consumption. Next, we construct a static MAC function that is fitted to be identical at each point in time (i.e., at each level of abatement on the optimal path of the exogenous TC model). This static MAC function will be concave, since TC makes the MAC function, which is assumed to be linear, fall over time. We fit based on a polynomial such that

$$MAC^{\%*} = \left(\sum_{n=0}^{N} \varphi_n^{nolearn} a_n^* \right), \tag{122}$$

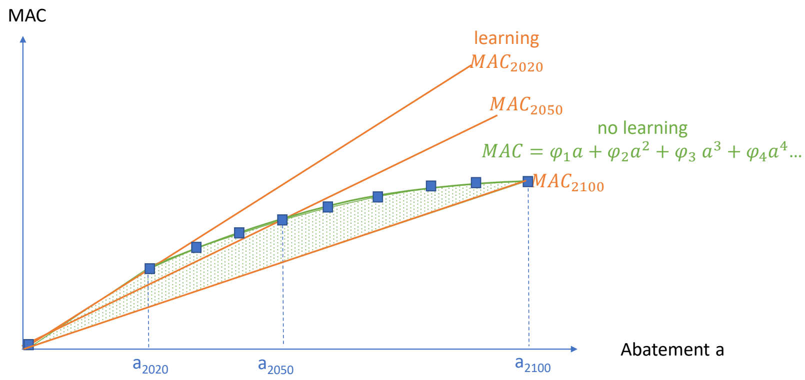

Figure A2: Graphical representation of the comparison between TC and no TC with identical but static MAC. The model with TC has a linear abatement cost function but with a time-dependent slope $\varphi_t a = \left(\varphi_\infty + (\varphi_0 - \varphi_\infty)e^{-g\varphi t}\right) a$ in the case of exogenous TC. This model results in an optimal set of $\{MAC^{\%*}, a^*\}$ depicted by the blue squares. Next, a polynomial in a (of order 15) is fitted to obtain a time-independent MAC that is identical on the optimal path. The polynomial also goes through the origin, in order to ensure similar total abatement costs, represented by the area under the curve. The dotted area is the difference between the total abatement costs of both models in 2100.

where the coefficients $\varphi_n^{nolearn}$ are independent of time, unlike in the model with exogenous TC.

We optimise the model with the static MAC function and check the solution is close to the optimum of the model with exogenous TC. Figure A3 shows that the correspondence is almost exact (less than 0.1% difference on emissions or MACs).

In fact, under certain specific functional forms the correspondence is exact in theory. In particular, if marginal abatement costs and marginal damages are proportional to consumption to the power ν and marginal utility is CES, the static MAC function is a *perfect* substitute for exogenous TC. The assumption of constant elasticities implies that we can factorise the marginal cost functions into a factor that does not depend on consumption and a power function of consumption, respectively: $c_a = -MAC^{\%}(a,t)c^\nu$; $c_T = -MD^{\%}(T)c^\nu$. Integrating Equation (45) yields

$$\lambda_t = \int_t^\infty e^{-(\delta-n)(u-t)} u_{c_u} \zeta c_u^\nu MD^{\%}(T) du. \tag{123}$$

Substituting Equation (44) gives

$$u_{c_t} c_t^\nu MAC^{\%}(a,t) = \int_t^\infty e^{-(\delta-n)(u-t)} u_{c_u} \zeta c_u^\nu MD^{\%}(T) du. \tag{124}$$

$u_{c_t} c_t^\nu$ is a constant (independent of time u) and can therefore be included in the integral,

$$MAC^{\%}(a,t) = \int_t^\infty e^{-(\delta-n)(u-t)} \left(\frac{c_u}{c_t}\right)^{\nu-\eta} \zeta MD^{\%}(T) du. \tag{125}$$

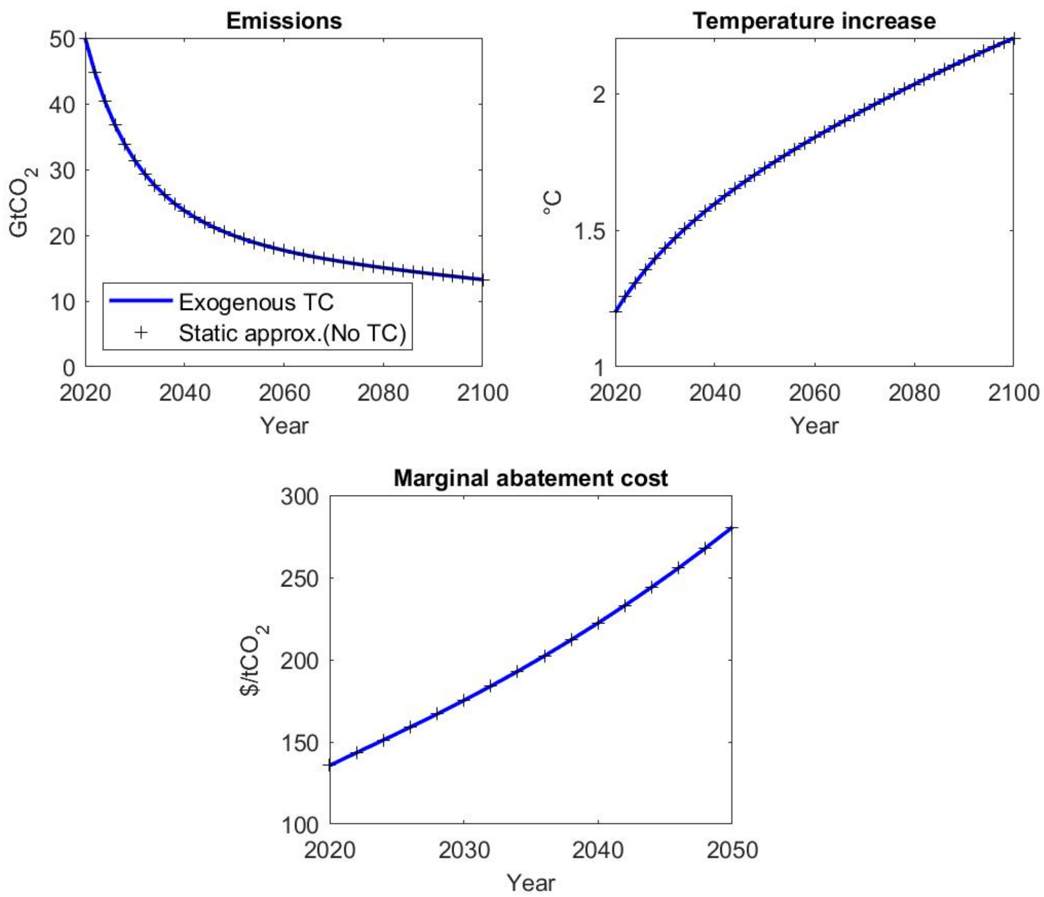

Figure A3: Exogenous TC versus no TC: fitting a polynomial in abatement to the MAC of the exogenous TC curve. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC on the bottom.

Call a^* the optimal abatement path of the model with exogenous TC. We try this path as a candidate solution to the model without TC. By assumption of identical MACs, the abatement path a^* will result in the same left-hand side of Equation (125) at each point in time. Given that we use the same climate model (Equation 40), the model without TC will result in the same temperature path. Figure A2 shows that although MACs are identical, total abatement costs will be higher in the model without TC (dotted area). Hence, consumption will be slightly lower in the model without TC. But since we assume $\nu = \eta$, the right-hand side of Equation 125 will also be identical at each point in time. Hence our candidate solution solves the Euler equation and is the optimal solution of the model without TC.

What if $\eta > \nu$? Figure A2 shows that although MACs are identical, total abatement costs will be lower in the model with TC (dotted area). Hence, consumption will be slightly higher in the model with TC. As a result, the discount factor $\left(\frac{c_u}{c_t}\right)^{\varsigma-\eta}$ will be slightly lower, and MACs and abatement will be lower too. The effect is too small to be visible on a graph though.

E.3 Approximation of R&D-based TC by exogenous TC

In this subsection we compare a model with exogenous TC with a model with endogenous R&D which is unaffected by deployment.

Consider the general model of Appendix A, where the knowledge stock is independent of abatement, but responds to investment I in early-stage R&D, that is, $\dot{H} = \psi(t, H, I)$. Assume that this endogenous R&D is replaced by an exogenously decreasing MAC function, such that on that the optimised path, the MAC function is identical at each point in time. Then, the optimal emission and MAC paths are quasi-identical, they differ only to the extent that R&D investment costs reduce consumption.

To show this, ignore investment costs c_I for a moment. Call $H^*(t)$ and $I^*(t)$ the optimal knowledge stock and investment respectively of the model with early-stage R&D. Compare this to a model with exogenous TC, where the paths of the knowledge stock and investment respectively are replaced by an exogenous function of time $f(t)$, such that $\forall t : c\,(a, H^*(t), I^*(t), T, t) = c\,(a, f(t), T, t)$. The production function (which equals consumption in this setup) will be identical over the path, resulting in the same MAC and marginal damage functions. Since $\psi_a = 0$ in both cases, the optimal abatement path a^* of both models satisfy Equation (49) for cost-benefit analysis and Equation (57) for cost-effectiveness analysis. Hence the exogenous TC model and the early-stage R&D model have the same optimal abatement path.

Now bring back the R&D investment costs – these reduce consumption and therefore alter marginal damages and the discount rate. The critical observation is that this effect on consumption growth will generally be negligible, given the small size of the investment costs when converted into a growth impact (IPCC, 2022). This is basis of Remark 4 in the main text.

Two caveats are in place though. First, note that although the technology function ψ is independent of abatement, the dependence on investment makes TC endogenous, that is, investment in green technology becomes more attractive in a lower emissions scenario. This is also why solving the model for the optimal R&D investment rule is not straightforward, as investment depends on the abatement and *vice versa*.

Second, whenever policy is not optimal, or whenever beliefs on future costs and damages are updated, the two models will no longer behave in the same way, because exogenous TC is independent of policy whereas investment in R&D depends on forward-looking policy.

Appendix F GMM results by model

	No TC	Exogenous TC					Endogenous TC				
	All models	All models	Exog models	FE: initial MAC	FE: initial MAC & quadr	FE: initial MAC & TC	All models	Endog models	FE: initial MAC	FE: initial MAC & quadr	FE: initial MAC & TC
θ	1.5e-03***	-1.7e-05	-2.9e-05	7.7e-04***	idem	4.7e-04***	5.1e-05	9.2e-05	6.6e-04***	idem	4.5e-04***
φ	2.7e-05***	1.2e-04***	1.2e-04***				1.4e-04***	1.3e-04***			
g_φ		.048***	.049***	.013***	.011***						
φ_∞		1.9e-05***	1.5e-05***	2.2e-07	idem	1.7e-05***					
φ_1				6.9e-05***	6.8e-05***	1.3e-04***			8.2e-05***	7.1e-05***	2.2e-04***
φ_2				1.3e-04***	1.2e-04***	9.7e-05***			1.5e-04***	1.2e-04***	8.0e-05**
φ_3				1.1e-04***	1.0e-04***	7.7e-05***			1.3e-04***	1.1e-04***	9.9e-05***
φ_4				4.3e-05***	4.6e-05***	8.7e-05***			5.1e-05***	5.0e-05***	8.7e-05***
φ_5				4.5e-05***	4.6e-05***	1.4e-04***			5.1e-05***	4.9e-05***	9.1e-05***
φ_6				5.3e-05***	5.6e-05***	8.6e-05***			6.4e-05***	6.1e-05***	9.1e-05***
φ_7				7.2e-05***	7.2e-05***	8.0e-05***			8.8e-05***	7.8e-05***	1.0e-04***
φ_8				1.8e-04***	1.5e-04***	1.7e-04***			2.1e-04***	1.5e-04***	2.0e-04***
φ_9				9.3e-05***	6.6e-05***	1.3e-04***			1.1e-04***	6.9e-05***	1.2e-04***
φ_{10}				2.0e-05***	2.6e-05***	8.3e-06***			2.3e-05***	3.1e-05***	3.9e-06***
φ_{11}				1.9e-04***	1.7e-04***	1.8e-04***			2.3e-04***	1.9e-04***	1.3e-04***
φ_{sq}					-4.0e-07***					-2.8e-07***	
$g_{\varphi 1}$						.038***					
$g_{\varphi 2}$						8.8e-03					
$g_{\varphi 3}$						9.5e-03***					
$g_{\varphi 4}$						.056***					
$g_{\varphi 5}$						.068***					
$g_{\varphi 6}$						.041***					
$g_{\varphi 7}$						.023***					
$g_{\varphi 8}$						.013***					
$g_{\varphi 9}$						.026***					
$g_{\varphi 10}$						3.0e-03**					
$g_{\varphi 11}$						.011***					
χ							.361***	.302***	.315***	.216***	
A_0							16.3***	12.7***	92.5***	idem	36.6***
χ_1											.515***
χ_2											.058
χ_3											.164***
χ_4											.387***
χ_5											.402***
χ_6											.336***
χ_7											.274***
χ_8											.213***
χ_9											.267***
χ_{10}											-.236***
χ_{11}											-.064
N	12393	12393	6545	12393	12393	12393	12393	5848	12393	12393	12393
ll	4.9e+04	5.0e+04	2.7e+04	5.2e+04	5.3e+04	5.2e+04	5.0e+04	2.3e+04	5.2e+04	5.2e+04	5.2e+04

Table A1: Parameter estimates for fitting both total and marginal abatement costs to the IPCC and NGFS databases of IAM results, including model-specific estimates. Models are grouped as follows: (1) AIM, (2) C3AIM, MERGE or C_ROADS, (3) GCAM, (4) GEM, (5) MESSAGE, (6) REMIND, (7) WITCH, (8) POLES, (9) IMAGE, (10) COFFEE, (11) DNE. Idem indicates that the value of the preceding (left) model is imposed. *** indicates significance at the 1% level.

Appendix G Additional quantitative results

Figure A4 presents optimal, cost-benefit climate policies in the absence of capital inertia (i.e., no abatement speed penalty). Without inertia, the social planner is free to choose large variations in initial emissions. Accordingly, initial emissions are much lower than in the presence of capital inertia, regardless of the existence and type of TC. This results in more slowly increasing temperatures, but since capital inertia is relevant in the short to medium run but less so in the long run, temperatures in 2100 are similar.

Consistent with the theoretical results in Appendix B, emissions under exogenous TC are initially higher (abatement is lower) than without TC, with a lower MAC/carbon price. Conversely, under endogenous, learning-based TC emissions are initially lower than without TC, with a higher MAC. The abatement path is steeper under both forms of TC, which is also consistent with the theory.

Figures A5 to A8 present optimal, cost-benefit climate policies with low/high discount rates and slow/fast TC. These plots provide further details on the results summarised in Table 4.

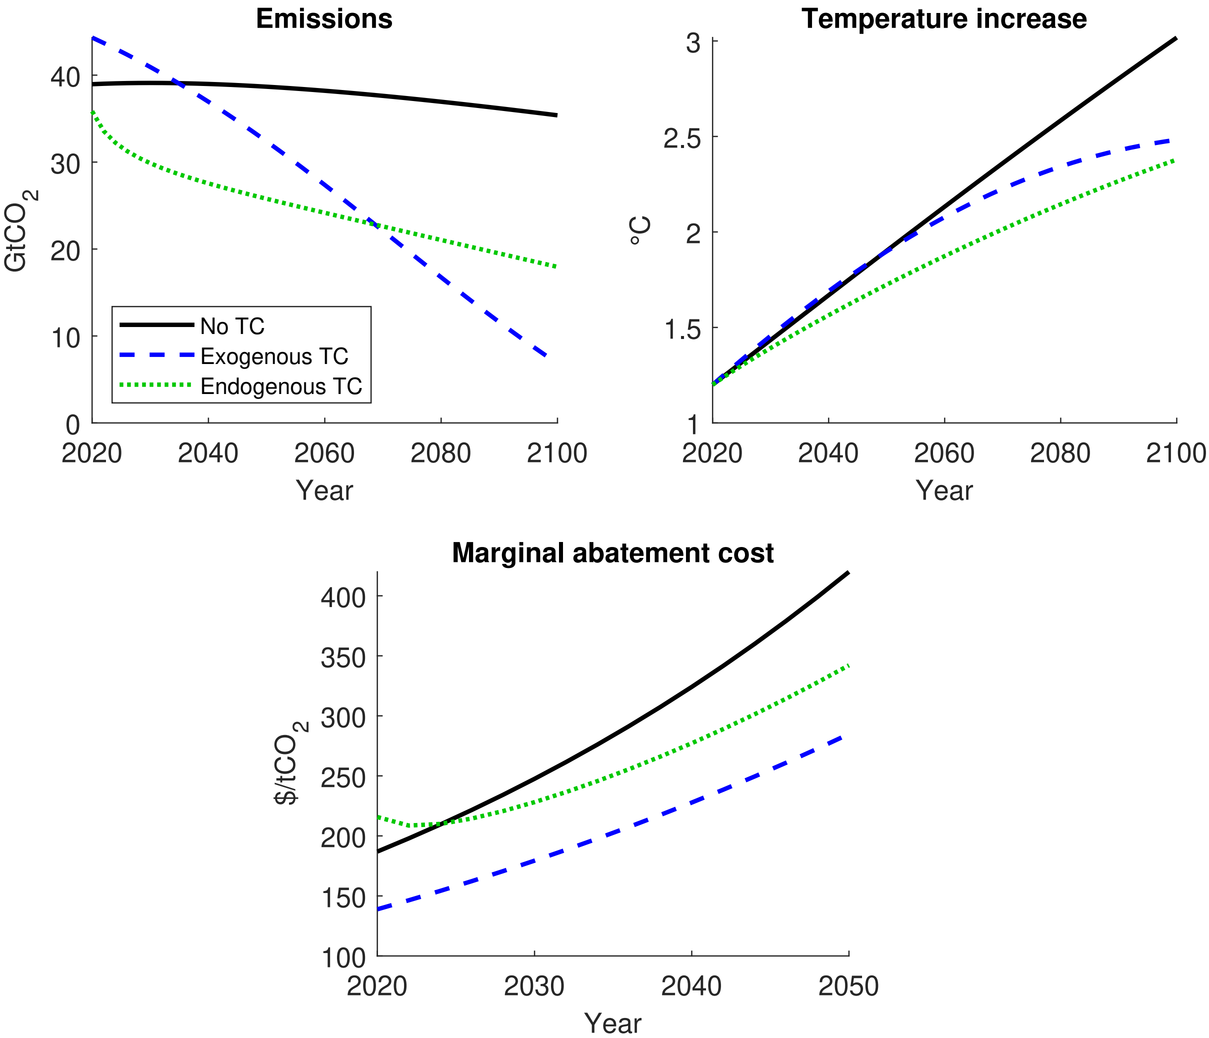

Figure A4: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC, for the case of **no capital inertia**. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom. Note that the modest increase in emissions in scenarios without TC is the result of population growth and marginal damages, which both create an incentive for early abatement (both factors exceed the effect of the discount rate $r + g - 2\frac{\dot{\sigma}}{\sigma}$ in Equation (78)).

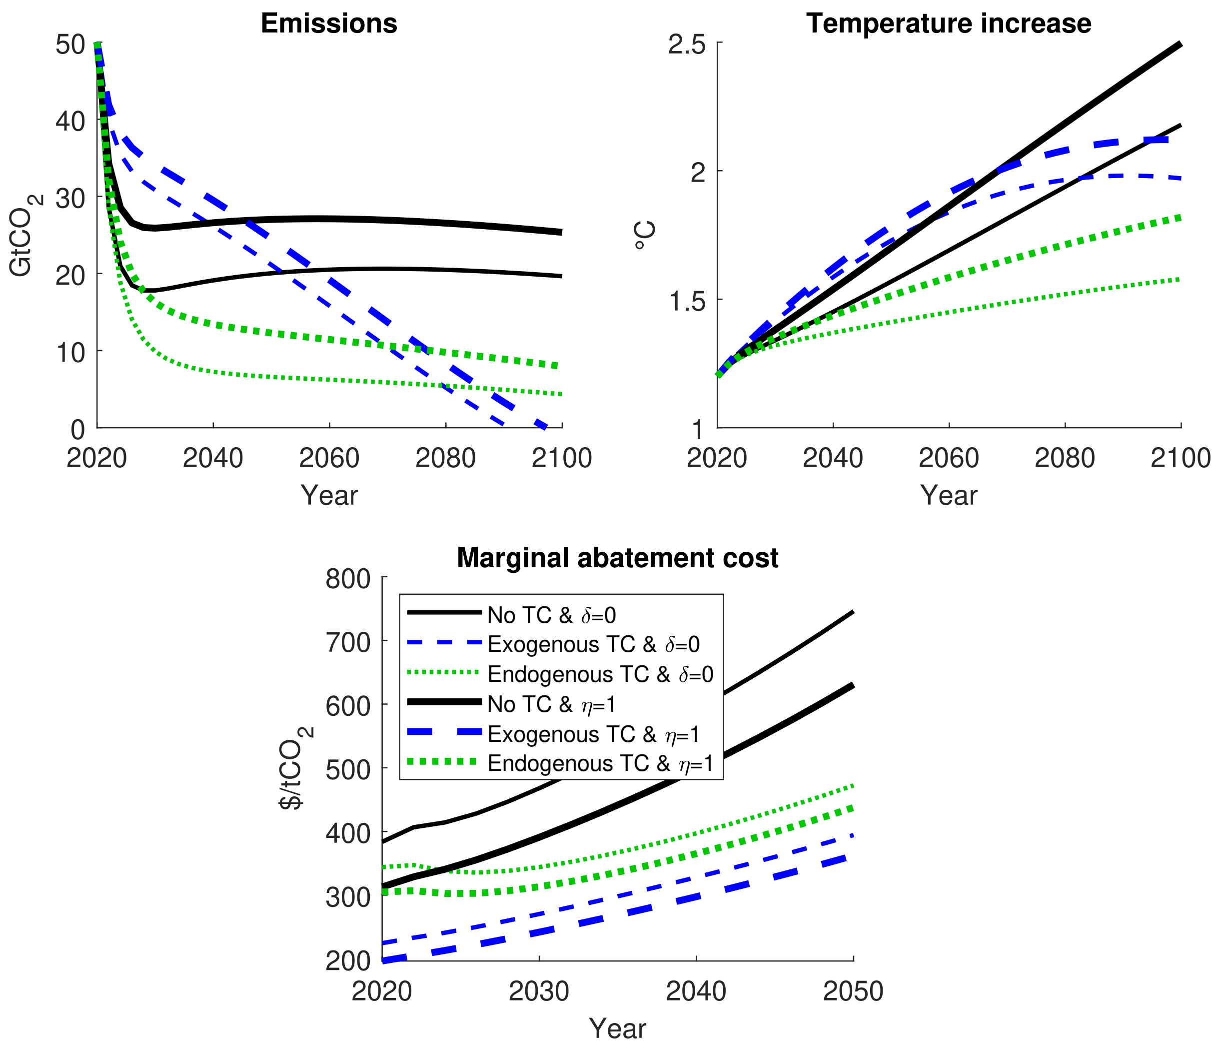

Figure A5: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC, under **low discount rates** ($\delta = 0$ or $\eta = 1$). Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

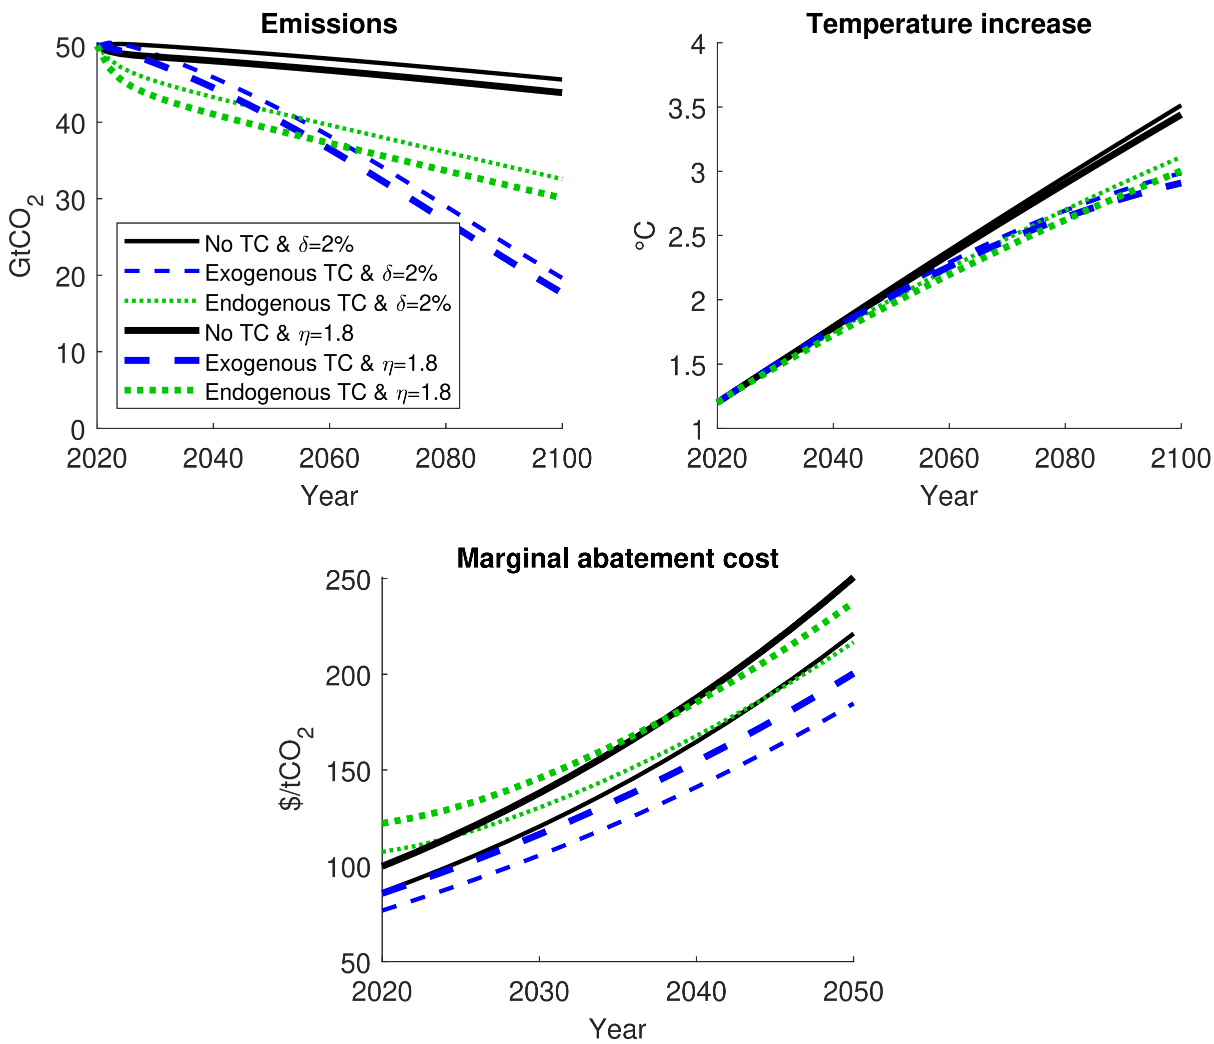

Figure A6: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC, under **high discount rates** ($\delta = 0.02$ or $\eta = 2$). Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

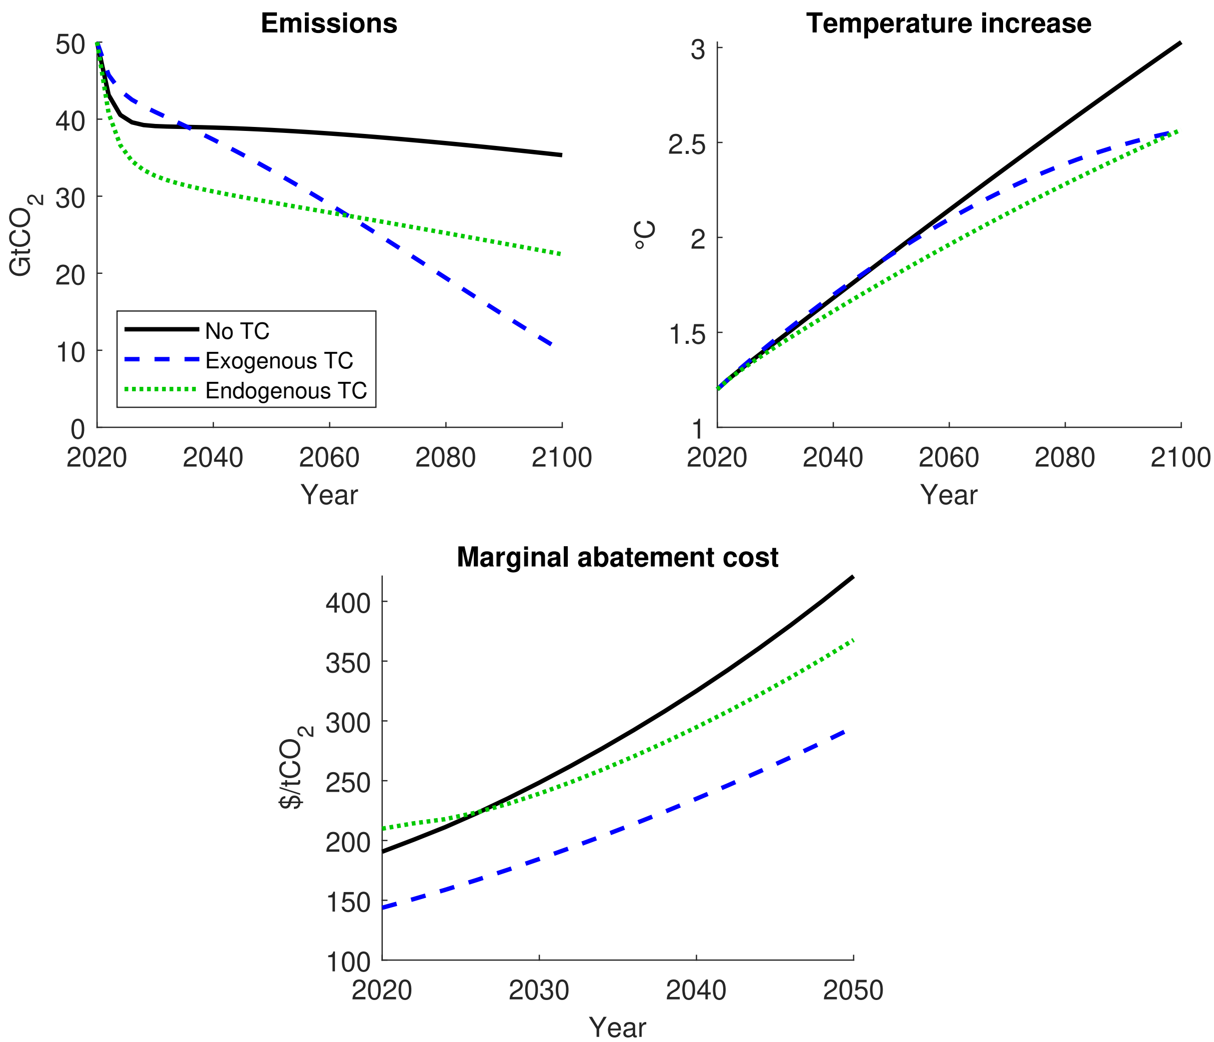

Figure A7: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC, under **slow TC**. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.

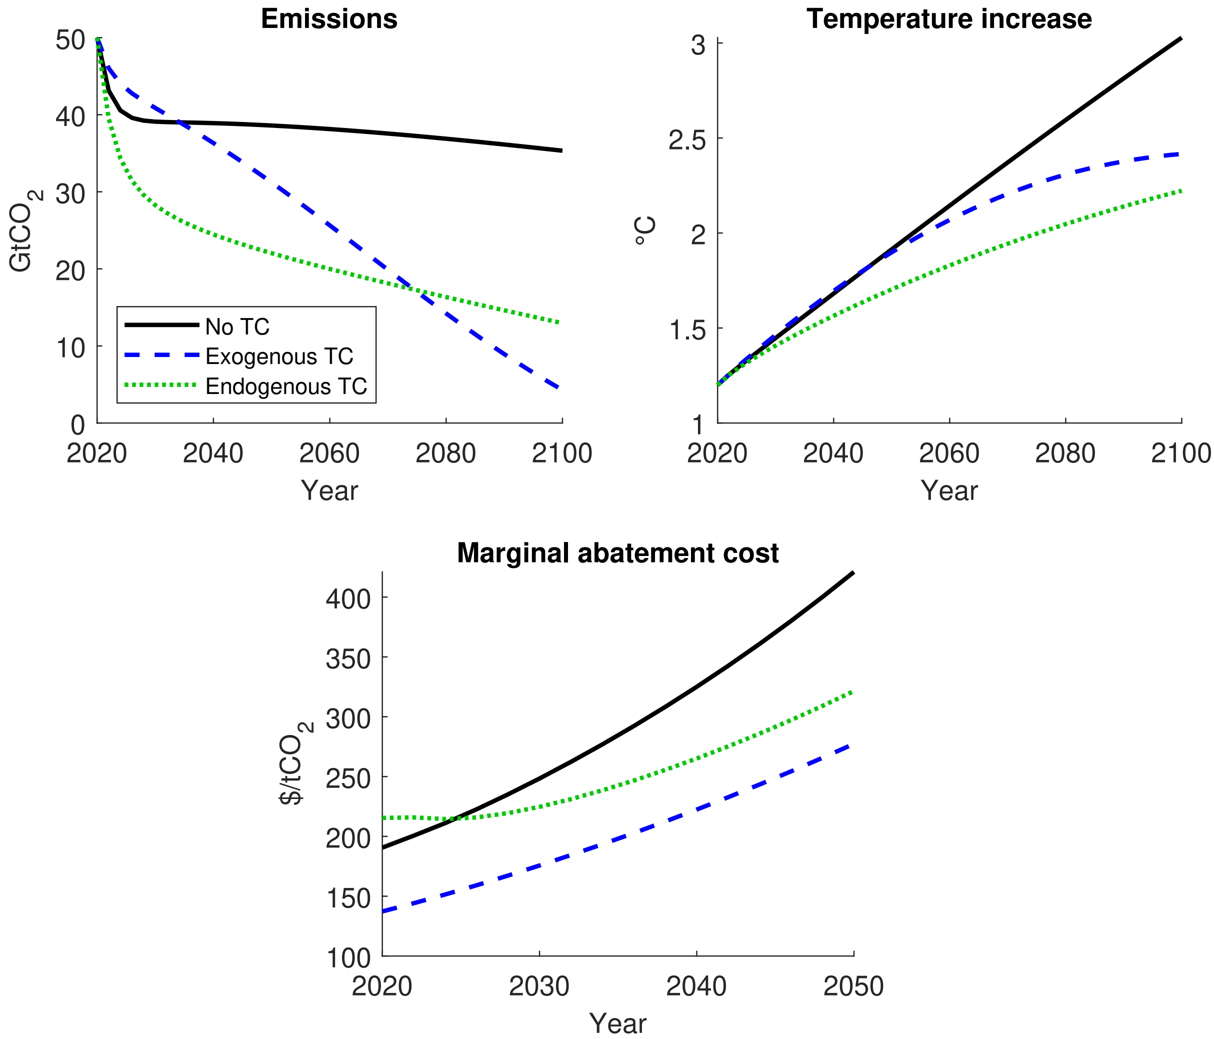

Figure A8: Optimal, cost-benefit climate policies without TC, with exogenous TC and with endogenous TC, under **fast TC**. Emissions in the top-left panel, temperature above pre-industrial in the top right, and the MAC/carbon price on the bottom.