DIW BERLIN

2161

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2026

Meeting Climate Targets under Distortionary Fiscal Policy: Directed Technical Change and Learning-by-Doing

Sonja Dobkowitz

IMPRESSUM

DIW Berlin
German Institute for Economic Research
Anton-Wilhelm-Amo-Str. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
http://www.diw.de

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
http://www.diw.de/discussionpapers

Discussion Papers of DIW Berlin are indexed in RePEc and Econstor:
http://ideas.repec.org/s/diw/diwwpp.html
https://www.econstor.eu/handle/10419/10

Meeting Climate Targets under Distortionary Fiscal Policy: Directed Technical Change and Learning-by-Doing

Sonja Dobkowitz *

April 24, 2026

Abstract I study optimal implementation of climate targets in a model with distortionary fiscal policy, learning-by-doing, and directed technical change. The key mechanism is that fiscal constraints link innovation policy to labor allocation, creating a tension between directing research and directing learning-by-doing. Analytically, I show that learning-by-doing shapes the effectiveness of carbon taxation in directing research through an expertise effect: carbon taxes are more effective at steering innovation toward green technologies when green expertise is relatively high. Quantitatively, I calibrate the model to the U.S. economy to characterize the optimal policy mix consistent with climate targets. I find that carbon should be taxed heavily, persistently exceeding the social cost of carbon. While higher carbon prices raise green expertise, they induce an excessively rapid reallocation of researchers from fossil to green technologies, generating persistent innovation misallocation. A welfare analysis shows that learning-by-doing substantially amplifies the cost of distortionary taxation, in particular during the transition to net-zero emissions.

JEL classification: H21, H23, O38, Q54, Q55

Keywords: Second-best climate policy, Directed technical change, Learning-by-doing, Ramsey taxation, Misallocation of innovation, Emissions target implementation

* DIW Berlin, sdobkowitz@diw.de, https://sonjadobkowitz.wordpress.com/.
I am deeply grateful to my supervisors Keith Kuester, Pavel Brendler, and Farzad Saidi for their support and guidance. I would like to thank Philippe Aghion, Tobias Bergmann, Janosch Brenzel-Weiss, Rubén Domínguez-Díaz, Beatriz Gaitan, Anna-Maria Goeth, Rob Hart, Thomas Hintermaier, Peter Kjær Kruse-Andersen, Thomas Köhler, Antonia Kurz, Kai Lessmann, Leonie Meissner, and Paul Schäfer for helpful discussions and comments. I further thank participants at the EEA Congress 2025, the EAERE Annual Conference 2025, VfS Annual Conference 2024, the EUI FSR Climate Annual Conference 2023, the PSE Summer School on Climate Change 2022, and the RTG-2281 Retreat 2022. Financial support by the German Research Foundation (DFG) through RTG-2281 "The Macroeconomics of Inequality" is gratefully acknowledged.

1 Introduction

The transition to net-zero emissions requires a fundamental reallocation of production. Governments rely on research subsidies to support this shift,[1] and a growing literature shows that such instruments facilitate the transition.[2] Yet financing subsidies requires distortionary taxation. At the same time, the transition depends on the accumulation of sector-specific expertise, driven by learning-by-doing.[3] Because distortionary taxes reduce labor supply, they hinder the accumulation of expertise that is critical for adopting green technologies. Despite their central role, distortionary fiscal policy is largely absent from the literature on both directed technical change and learning-driven green transitions.[4] This paper studies optimal implementation of emissions targets when learning and innovation are jointly constrained by distortionary fiscal policy.

The main result is that to steer learning-by-doing during the transition, the planner relies heavily on carbon taxation, accepting exacerbated labor supply distortions as a consequence. As a result, second-best carbon taxes persistently exceed the social cost of carbon.[5] While this policy accelerates the accumulation of green expertise, it induces an excessively rapid reallocation of research from fossil to green technologies. Because

[1]Under its Horizon Europe program, for instance, the European Union has allocated 35% of its grants to research aligned with the European Green Deal (European Commission, 2020). In the U.S., the Inflation Reduction Act of 2022 earmarks approximately $369 billion for energy security and climate change initiatives, including funding for clean energy R&D (United Nations Conference on Trade and Development, 2022).

[2]On the theoretical side, Acemoglu et al. (2012) show that research subsidies optimally complement carbon taxes in the green transition by addressing dynamic knowledge spillovers. Quantitative studies further explore their role in directed technical change frameworks Acemoglu et al. (2016); Hart (2019). On the empirical side, Howell (2017) shows that government R&D subsidies in the energy sectors are especially beneficial for financially constrained firms, highlighting an important role for clean energy technologies. Other contributions include Johnstone et al. (2010) and Popp (2016).

[3]Barwick et al. (2025) find a substantial role of learning-by-doing in reducing the cost of electric vehicles. Irwin and Klenow (1994) and Nykvist and Nilsson (2015) document similar effects, while smaller effects are reported by Bollinger and Gillingham (2023) and Nemet (2006).

[4]In the directed technical change literature, influential contributions such as Acemoglu et al. (2012, 2016); Hart (2019) characterize optimal combinations of carbon pricing and research subsidies under the assumption of non-distortionary financing. Related work on learning-by-doing in the green transition, such as Goulder and Mathai (2000); Kverndokk and Rosendahl (2007); Rezai and van der Ploeg (2017), similarly abstracts from fiscal distortions. See the literature review below for a more detailed delineation.

[5]The social cost of carbon is defined by the shadow price of the emissions constraint in the Ramsey problem divided by the household's marginal utility of consumption. In the absence of explicit climate damages, this object captures the scarcity value of the emissions constraint rather than marginal environmental damages. The definition follows Belfiori and Rezai (2024), with direct utility and production damages of emissions set to zero. In a first-best environment with lump-sum taxation and instruments that directly target sector-specific learning, the optimal carbon tax coincides with this object.

research subsidies are financed through distortionary taxation, the planner only partially offsets these incentives, tolerating a persistent misallocation of innovation. Together, these results highlight a policy trade-off between directing learning to the green sector and supporting a smooth transition of research toward low-carbon technologies.

In more detail, I develop a dynamic model that integrates three core features of the green transition: learning-by-doing, directed technical change, and elastic labor supply. In the model, energy from fossil sources emits CO_2 but can be partially substituted by low-carbon, *green* alternatives in production. The government chooses time paths for distortionary labor income taxes, carbon taxes, and research subsidies to maximize welfare, subject to an exogenous budget constraint.[6] It anticipates that net emissions must be limited in the near term and eliminated entirely in the long run. Calibrating the model to the U.S. economy, I characterize the optimal policy mix during the transition toward net-zero emissions under a 2°C-consistent dynamic emissions target.[7]

Before turning to the quantitative analysis, I use the model to establish a general-equilibrium link between sectoral expertise and the allocation of research: research tends to flow toward sectors with more experienced workers, where new machines can be adopted more easily. This expertise effect shapes how strongly carbon taxes shift innovative activity toward the green sector. Under plausible parameter values, a marginal increase in the carbon tax is more effective, all else equal, when workers already have greater experience with green technologies. In addition, carbon taxation influences the allocation of researchers indirectly by increasing green expertise as production expands.

[6] I abstract from targeted subsidies to foster sector-specific expertise, treating learning as an endogenous outcome of production and labor supply. This approach follows the public finance literature on optimal taxation with learning-by-doing (e.g. Peterman, 2016; Makris and Pavan, 2021), which emphasizes the interaction between learning-by-doing and optimal fiscal tax design. Moreover, targeted learning subsidies would not eliminate the fiscal trade-offs at the core of the analysis, as they, as well, require distortionary financing.

[7] The *optimal policy* concept in the present paper refers to policies that maximize social welfare conditional on an exogenously specified emissions pathway. The analysis is therefore one of target-conditional optimal policy rather than the joint determination of optimal emissions and policy. In so doing, my paper contributes a cost-effectiveness analysis in a directed technical change setting. Most of the literature on environmental policies in directed technical change frameworks investigates environmental policies in a cost-benefit specification where the environmental externality enters through the utility function or production damages. Fried (2018) studies the implementation of a one-period static emission target, however, she abstracts from welfare maximization. Other papers studying cost-effectiveness trade-offs are, for instance, Goulder and Mathai (2000) and Edenhofer et al. (2005).

Together, these channels imply that carbon taxes stimulate green research more strongly and may substitute for research subsidies that require the use of distortionary fiscal instruments. Whether this mechanism helps mitigate the tension between innovation and learning, however, depends on how well the induced shift toward green research aligns with the planner's optimal innovation trajectory.

To discipline the quantitative model, I construct a dataset of patents related to renewable and fossil-based energy-supply technologies in the U.S. from 1950 to 2014. The classification follows an energy-technology taxonomy developed by the International Energy Agency and the European Patent Office. My implementation matches micro-level patents to technology types through Cooperative Patent Classification codes, complemented by a text-based search to ensure accurate identification of patents by energy source. This approach contributes a transparent and replicable method for measuring patents in renewable and fossil-related energy-supply technologies to the literature (e.g. Popp, 2002; Acemoglu et al., 2016). The data reveal that the stock of fossil energy technologies held a 25% lead over green alternatives at the outset of the transition in the 2010–2014 period.

I contribute three quantitative results. First, I find that learning-by-doing motivates the planner to rely more heavily on carbon taxation and less on labor income taxation. This result is noteworthy because carbon taxes are typically viewed as particularly distortionary instruments in settings with pre-existing distortionary taxes.[8] By contrast, public finance analyses emphasize that learning-by-doing can further raise the distortions of fiscal taxes by affecting both labor supply and future productivity (Peterman, 2016). From this perspective, one would expect the second-best policy to substitute away from carbon taxation toward labor income taxation.

Instead, I find that the optimal carbon tax remains more than 2% above the social cost of carbon throughout the transition. By shifting production toward green technologies,

[8] Starting from Bovenberg and De Mooij (1994), a large literature studies environmental policy under distortionary taxation; see, among others, Goulder (1995); Bovenberg and Goulder (1996); Metcalf (2003); Jacobs and van der Ploeg (2019); Barrage (2020); Douenne et al. (2026). The literature posits that, in general equilibrium, carbon taxes tend to be more distortionary than labor income taxes because they not only shrink their own tax base but also erode the labor-tax base by lowering real wages and labor supply.

carbon pricing accelerates the accumulation of green expertise. During the transition, this directional effect is quantitatively central: what matters for welfare is not only how quickly expertise accumulates, but where it is built. For this reason, the planner is willing to tolerate a more distortionary fiscal mix. Learning-by-doing thus simultaneously increases the distortionary cost of carbon taxation through its interaction with labor supply and strengthens the case for using it.

Second, I document a persistent misallocation of research effort under distortionary fiscal policy. Because research subsidies must be funded through distortionary taxes, the planner limits their use and tolerates a composition of research that deviates from a regime with lump-sum–financed subsidies.[9] In the implemented allocation, innovative activity transitions too rapidly to green technologies.[10]

I decompose the effect of endogenous learning-by-doing on the allocation of researchers into policy and expertise effects relative to an environment with exogenous expertise accumulation.[11] The higher carbon pricing set to support learning-by-doing directly reduces the fossil-to-green research ratio by between 2.5 and 2.8% during the transition, ceteris paribus. While higher subsidies to fossil research could, in principle, offset this induced shift, fiscal distortions limit how strongly subsidies can respond, even more so when learning-by-doing is endogenous. In the constrained regime, therefore, the subsidy response is weaker than under lump-sum finance, widening the misallocation of innovation.

Third, I quantify the welfare costs of fiscal distortions using consumption-equivalent variation. Endogenizing learning-by-doing raises these costs by about 15%, an amplification roughly twice as large as in a scenario without an emissions target.

[9] I define *misallocation* as the deviation of the implemented allocation from the allocation the planner would choose if research subsidies could be financed through lump-sum taxation. I refer to this latter allocation as the *target* allocation of scientists. A direct comparison to the non-distortionary benchmark would be misleading, because fiscal distortions themselves alter the target trajectory of innovation: the emissions target is implemented at a smaller scale, allowing for a higher share of fossil energy in production which in turn shapes the beneficial allocation of research.

[10] The effect is driven by decreasing returns to research through the risk to duplicate ideas and cross-sectoral knowledge spillovers. Barbieri et al. (2023) provide empirical evidence on knowledge spillovers from non-green to green technologies. An example are rotors in steam turbines used for electricity generation from fossil fuels, a technology that is closely related to technologies used in wind-energy generation.

[11] The benchmark fixes sectoral expertise at the path induced by the optimal policy under lump-sum taxation.

These welfare costs corroborate the mechanisms analyzed throughout the paper. Absent an emissions target, fiscal distortions are amplified because lower labor supply reduces future productivity through slower accumulation of expertise. Under an emissions target, endogenous learning-by-doing creates a trade-off between maintaining labor supply and using distortionary fiscal instruments to steer learning toward the green sector.

The amplification of fiscal distortions through learning-by-doing under an emissions target is driven mainly by distortions to learning. While learning-by-doing increases the welfare cost associated with innovation misallocation, this effect is quantitatively negligible in levels and does not affect the overall 15% amplification of fiscal distortions.

Literature This paper contributes to three strands of the literature. First, the paper adds to the mounting literature on climate policy in endogenous-growth models with directed technical change (e.g. Acemoglu et al., 2012; Hémous, 2016; Acemoglu et al., 2016; Van den Bijgaart, 2017; Greaker et al., 2018; Fried, 2018; Hart, 2019; Acemoglu et al., 2023; Lemoine, 2024; Aghion et al., 2025). In such settings, research subsidies should complement carbon pricing but are typically financed with lump-sum taxation. By introducing distortionary fiscal constraints, I revisit how research subsidies and carbon taxes should be combined to support the transition toward green technologies. Adding fiscal distortions means that carbon taxes and research subsidies compete for fiscal resources: while the latter relies on government funds, the former complicates raising them. The quantitative results highlight the importance of carbon taxes throughout the transition and in particular in the short run to boost workers' expertise on green technologies.

Moreover, I show that fiscal constraints change the target allocation of researchers. As the emission target is implemented at a lower scale and higher share of fossil energy, the target allocation of research efforts consists of a higher maintained share of fossil-related R&D relative to a non-distortionary fiscal setting. This mechanism is absent from existing directed innovation models, where fiscal finance does not constrain policy design.[12]

[12] For related work on environmental policies in general endogenous growth models and the combination of corrective taxes and research subsidies, see, among others, Bovenberg and Smulders (1995, 1996); Schneider and Goulder (1997); Goulder and Mathai (2000); Van der Zwaan et al. (2002); Popp (2004,

A second, contribution to the directed technical change literature is to show how learning-by-doing shapes the direction of research through a market-based mechanism. While the idea that learning-by-doing can generate self-reinforcing dynamics and technological lock-in is well established in the literature on technology adoption (e.g. Arthur, 1989; David, 1985; Jovanovic and Nyarko, 1996) this paper formalizes how such dynamics endogenously redirect the allocation of research effort across sectors in a directed innovation framework. In particular, it builds on Young (1993), who embeds learning-by-doing in a love-for-variety growth model where aggregate research responds to learning via increased profitability from cheaper production. The present paper extends this logic to the direction of innovation: research is endogenously directed toward the sector with greater accumulated experience. In the context of climate policy, the mechanism shapes how strongly research responds to carbon taxation.

Second, my paper adds to the literature on the role of learning-by-doing for environmental policies. These contributions study the role of technology subsidies and spotlight technological lock-in (e.g. Goulder and Mathai, 2000; Gerlagh et al., 2004; Kverndokk and Rosendahl, 2007; Kalkuhl et al., 2012, 2013; Rezai and van der Ploeg, 2017). Research-based innovation and learning-by-doing are almost always analyzed in isolation with few exceptions such as Edenhofer et al. (2005) and Fischer and Newell (2008). Goulder and Mathai (2000) compare the implications of learning-by-doing and R&D-based innovation for environmental policy design, pointing to differences depending on the source of growth. I show that studying learning-by-doing and directed technical change jointly entails relevant trade-offs and policy implications in distortionary fiscal settings—a dimension absent from the aforementioned studies.

Third, this paper contributes to the literature on optimal environmental policy in distortionary fiscal settings. A large body of work has examined how environmental taxation interacts with distortionary fiscal systems, often focusing on the fiscal implications of carbon pricing (e.g. Bovenberg and De Mooij, 1994; Goulder, 1995; Bovenberg and Goulder, 1996; Fullerton, 1997; Metcalf, 2003; Jacobs and van der Ploeg, 2019; Douenne

2006); Fischer and Newell (2008); Gerlagh et al. (2009, 2014). All these papers abstract from fiscal constraints.

et al., 2026). Albeit abstracting from endogenous growth, Barrage (2020) is closely related by studying optimal environmental taxation in a dynamic framework with fiscal distortions. I discuss the insights from this literature in an induced technical change setup. In particular, my findings point to the additional gains of carbon taxes that outweigh the fiscal costs which commonly lower their level below the SCC.

A smaller set of papers embeds corrective taxation in endogenous growth models with distortionary fiscal regimes, showing that environmental policy can foster long-run growth, for example when tax burdens are shifted toward less elastic tax bases or when environmental quality enters production (Bovenberg and Mooij, 1997; Hettich, 1998; Greiner, 2005; Fullerton and Kim, 2008). My contribution is twofold. First, while these papers largely focus on long-run growth, I shift the perspective to the trajectory toward net-zero emissions. Second, rather than treating growth as a homogeneous process summarized by a balanced-growth rate, I place the analysis in a dynamic framework where growth is explicitly structured to meet an exogenous emissions target. From this perspective, the key question is not whether environmental policy promotes long-run growth, but how the structure of growth should be engineered to meet a predetermined emissions target.

Outline The remainder of the paper is structured as follows. Section 2 presents the model. I derive theoretical results in Section 3. I calibrate the model in Section 4 before Section 5 discusses the quantitative results. Section 6 concludes.

2 Model

This section presents the quantitative framework. I construct a model of directed technical change building on the frameworks developed in Acemoglu et al. (2012) and (Fried, 2018). I extend previous work in two essential ways. First, I allow for elastic labor supply. Second, I add learning-by-doing.

Households A representative household describes the household side. The household chooses consumption, C_t, and the share of hours spent working, H_t, taking prices as given. The household owns machine producing firms from which it receives profits, Π_t.[13] It also supplies scientific work in a fixed amount: S.[14] The household behaves according to solving below problem each period:

$$\max_{C_t, H_t} \quad \log(C_t) - \chi \frac{H_t^{1+\sigma}}{1+\sigma}$$
$$s.t. \quad p_t C_t \leq w_t(1-\tau_{lt})H_t + w_{st}S + Gov_t^T + \Pi_t - T_{xt}.$$

The variables w_t, w_{st}, and p_t indicate prices for labor, research, and the final consumption good. Lump-sum transfers from the carbon tax and the labor income tax minus subsidies for research are denoted by Gov_t^T. Labor income is taxed at a linear rate τ_{lt}. The utility specification follows King et al. (1988) admitting balanced growth with constant hours worked. The parameter σ captures the inverse Frisch elasticity and χ scales the disutility of labor. T_{xt} indicates lump-sum transfers to finance subsidies for the correction of monopolistic competition.[15]

Production Production separates into final good production, energy production, intermediate good production, and the production of machines. The final sector is perfectly competitive combining non-energy and energy goods according to:

$$Y_t = \left(\delta_y^{\frac{1}{\varepsilon_y}} E_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1-\delta_y)^{\frac{1}{\varepsilon_y}} N_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}} \right)^{\frac{\varepsilon_y}{\varepsilon_y - 1}}.$$

[13] Where $\Pi_t = \sum_J \left(\int_0^1 \pi_{xJit} \right) di$ is the sum over profits of all firms i in each sector J.

[14] It is common to fix the supply of scientists in the literature on directed technical change (Acemoglu et al., 2012; Fried, 2018). The assumption mutes the importance of the level of research and helps focus the discussion on the allocation of research which is the purpose of this paper. On the downside, it implies that increasing research in one sector translates into a crowding-out of research in other sectors (compare Hémous and Olsen, 2021).

[15] Allowing for such instruments enables me to abstract from distortions arising from the monopolistic structure in the machine producers' market. Therefore, also in the distortionary fiscal setting, I allow for lump-sum taxes to finance such subsidies.

I take the final good as the numeraire and define its price as $p_t = \left(\delta_y p_{Et}^{1-\varepsilon_y} + (1-\delta_y) p_{Nt}^{1-\varepsilon_y}\right)^{\frac{1}{1-\varepsilon_y}}$. Energy producers perfectly competitively combine fossil and green energy to a composite energy good:

$$E_t = \left(F_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}} + G_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}}\right)^{\frac{\varepsilon_e}{\varepsilon_e - 1}}.$$

The price of energy is determined as $p_{Et} = \left((p_{Ft} + \tau_{Ft})^{1-\varepsilon_e} + p_{Gt}^{1-\varepsilon_e}\right)^{\frac{1}{1-\varepsilon_e}}$. The government levies a sales tax per unit of fossil energy bought by energy producers, τ_{Ft}, which I refer to as *carbon tax*.

Intermediate goods, fossil, F_t, green, G_t, and non-energy, N_t, are produced in competitive sectors using labor and sector-specific machines. The production function in sector $J \in \{F, G, N\}$ reads

$$J_t = L_{Jt}^{1-\alpha_J} \int_0^1 A_{Jit}^{1-\alpha_J} x_{Jit}^{\alpha_J} di.$$

The variable A_{Jit} indicates the productivity of machine i in sector J at time t, x_{Jit}. Productivity in the non-energy sector evolves according to $A_{Nt} = (1+\gamma_n) A_{Nt-1}$. Intermediate good producers maximize profits:

$$\pi_{Jt} = p_{Jt} J_t - w_t L_{Jt} - \int_0^1 p_{xJit} x_{Jit} di,$$

where p_{xJit} denotes the price of machines from producer i in sector J. The wage rate for labor, w_t, is equal across sectors due to free labor mobility.

Machine producers are imperfect monopolists seeking to maximize profits. They choose the price at which to sell their machines to intermediate good producers. While machine producers in the non-energy sector take technology growth of their machines as given, and the amount of scientists is fixed $s_{Nit} = 0$ for all non-energy firms in all periods, machine producers in the fossil and green sector decide on the amount of scientists to employ. Demand for machines increases with their productivity which again is a function of technological progress. This provides the incentive to invest in research.

Irrespective of the sector, the costs of producing one machine is set to one unit of the final output good similar to Fried (2018) and Acemoglu et al. (2012). Following the same literature, machine producers only receive returns to innovation for one period. Afterwards, patents expire.[16] Taken together, machine producer i's profits in sector J are given by

$$\pi_{xJit} = p_{xJit}(1+\zeta_{Jt})x_{Jit} - x_{Jit} - w_{st}\left(1-\tau_{sJt}\right)s_{Jit}. \tag{1}$$

The government subsidizes machine production by ζ_{Jt} financed by lump-sum taxes on the household to correct for the monopolistic structure.

The government can subsidize green and fossil research via τ_{sJt}. In the absence of fiscal distortions, only the relative wedge between sectoral subsidies is determined by equilibrium conditions, making the normalization of one subsidy innocuous. In the presence of fiscal distortions, however, such normalization is no longer without loss of generality, as it implicitly determines which sector's subsidy entails distortionary financing. To ensure fiscal neutrality across sectors, I allow for two distinct, strictly positive subsidies, avoiding normalization that would arbitrarily favor one sector over the other.

Innovation Innovation drives the accumulation of knowledge and technological progress. In the model, it arises from research activity, where scientists' productivity depends on the existing stock of knowledge. I use the terms *knowledge* and *technology* to distinguish between innovation as an input to innovation and as an input to the production of the consumption good, respectively.

The law of motion for the knowledge stock of firm i in sector J is given by:

$$K_{Jit} = (1-\delta_K)K_{Jt-1} + \gamma\left(s_{Jit}\right)^{\eta}\left(K_{-J,t-1}^{\phi}K_{Jt-1}^{1-\phi}\right)^{1-\nu}, \tag{2}$$

where $K_{Jt} = \int_0^1 K_{Jit}\,di$. The parameter γ scales the productivity of researchers, and δ_K denotes the depreciation rate of knowledge. The parameter $\eta < 1$ introduces diminishing

[16] This choice is motivated by Fried (2018) who provides evidence for 5-years patents to match empirically well with knowledge spillovers.

returns to research, capturing a *stepping-on-toes* effect. As more researchers work on the same machine, the likelihood of duplicating ideas increases. A more equal allocation of researchers across sectors alleviates this risk.

Private incentives to invest in research diverge from the social optimum since accumulated knowledge raises the productivity of future innovation: a *building-on-the-shoulders-of-giants* effect. The model differentiates two types of spillovers. First, firms do not internalize the effect of today's research on future research productivity within their own sector, which is governed by the term $K_{Jt-1}^{(1-\phi)(1-\nu)}$. Second, they neglect cross-sectoral knowledge spillovers from innovation originating in the other sector, captured by $K_{-J,t-1}^{\phi(1-\nu)}$. This channel is essential to understand why the optimal policy seeks to maintain some fossil research activity during the transition to net-zero production.

The composite knowledge term, $K_{Jt-1}^{1-\phi}K_{-J,t-1}^{\phi}$, enters the innovation process with decreasing returns due to $\nu \in (0, 1)$. This *fishing-out* effect reflects the idea that early innovations are more important, while later ones tend to be incremental (Jones and Williams, 1998). As a result, maintaining a given rate of knowledge growth requires increasing research effort over time. Empirical evidence supports the presence of such effects and the associated slowdown in long-run growth (Bloom et al., 2020; Kruse-Andersen, 2023).[17]

Productivity, technology, and learning-by-doing The productivity of machines and technology are linked by how well workers can handle technologies, i.e., their *expertise*, Q_{Jt}.[18] Consistent with the free mobility of labor, workers' expertise is sector specific and evolves with sectoral output according to

$$Q_{Jt} = (1 - \delta_K)Q_{Jt-1} + J_t. \tag{3}$$

This specification has two advantages. First, it allows for a direct matching with empirical papers that study the effects of cumulative output on total factor productivity in the

[17] Fishing-out effects imply that for a balanced growth path (BGP) to exist, the number of scientists has to grow. In the calibration, I assume that on the BGP the supply of scientists grows at its BGP-consistent rate. In the base period, 2015-2019, it has reached the level of unity at which it is fixed thereafter.

[18] Throughout the paper, learning-by-doing refers to the mechanism through which production activity increases sector-specific expertise, while expertise denotes the corresponding state variable.

calibration. Second, the additive formulation admits a balanced growth path. Expertise decays at the same rate as technological knowledge accounting for the idea that experience in handling an outdated machine does not fully help employing a new technology (Young, 1993).

Expertise and technological knowledge jointly determine labor-augmenting productivity:

$$A_{Jit} = \kappa Q_{Jt}^{\iota_L} K_{Jit}^{\iota_K}. \tag{4}$$

The level shifter, κ, is common to all sectors. It allows to match observed output and initial knowledge stocks from the data by subsuming normalization in the calibration process. The parameter ι_K captures the elasticity of productivity to innovation, and ι_L analogously governs the elasticity with respect to expertise. Both elasticities are equal across sectors. This assumption together with sector-specific labor shares implies that the elasticity of total factor productivity, $A_{Jt}^{1-\alpha_J}$, to knowledge and expertise is sector specific and increases with the labor share.[19] I discuss the existence and stability of a balanced-growth path (BGP) in Appendix A.

Markets In equilibrium, markets clear. I explicitly model markets for workers, scientists, and the final consumption good:

$$\begin{aligned} H_t &= L_{Ft} + L_{Gt} + L_{Nt}, \\ S &= S_{Ft} + S_{Gt} \\ C_t &= Y_t - (X_{Ft} + X_{Gt} + X_{Nt}), \end{aligned}$$

where $S_{Jt} = \int_0^1 (s_{Jit})\, di$ and $X_{Jt} = \int_0^1 (x_{Jit})\, di$ which follows from the symmetry of machine producing firms within a sector. I assume free movement of scientists across sectors, which is justified by the 5-year duration of one period and research only occurring in the energy sectors which are closely related in terms of research fields.

[19]See Subsection E.3 for a detailed discussion.

Competitive Equilibrium Before turning to the Ramsey problem, I define a competitive equilibrium as

Definition 1. *Given initial values for the stocks of expertise, knowledge, and non-energy productivity* $\{Q_{F0}, Q_{G0}, K_{F0}, K_{G0}, A_{N0}\}$*, a competitive equilibrium characterizes the sequence of allocations* $\{C_t, H_t, L_{Ft}, L_{Gt}, L_{Nt}, s_{Fit}, s_{Git}, x_{Fit}, x_{Git}, x_{Nit}\}$*, prices* $\{w_{lt}, w_{st}, p_{Ft}, p_{Gt}, p_{Nt}, p_{xFit}, p_{xGit}, p_{xNit}\}$*, and a sequence of policies* $\{\tau_{Ft}, \tau_{lt}, \tau_{sFt}, \tau_{sGt}, \zeta_{Ft}, \zeta_{Ft}, \zeta_{Nt}\}$ *such that (i) households and firms optimize given prices and policies, (ii) the temporal government budget is balanced, and (iii) markets clear.*

Government The government seeks to maximize lifetime utility of the representative household. It is characterized as a Ramsey planner who chooses and commits to a sequence of policies in the initial period. The government is constrained by a dynamic emissions target, Ω_t, and has to balance its budget. It takes the behavior of firms and households as given discounting period utility with the household's time discount factor, β. The planner chooses time paths for carbon taxes, labor income taxes and research subsidies solving:

$$\max_{\{\tau_{lt}\}_{t=0}^{\infty}, \{\tau_{Ft}\}_{t=0}^{\infty}, \{\tau_{sFt}\}_{t=0}^{\infty}, \{\tau_{sGt}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left(\log(C_t) - \chi \frac{H_t^{1+\sigma}}{1+\sigma} \right)$$

$$s.t.\ \ \omega F_t - \delta \leq \Omega_t, \tag{5}$$

$$\tau_{Ft} F_t + \tau_{lt} w_t H_t - T_{Rt} = Gov_t^T, \tag{6}$$

$$Gov_t^T \geq Gov_t^{min} \tag{7}$$

subject to the behavior of firms and households, and feasibility.[20] Constraint (5) is the emissions target. The parameter δ captures the capacity of the environment to reduce emitted CO_2 through natural sinks, such as forests and moors. The parameter ω determines CO_2 emissions per unit of fossil energy produced. The budget, constraint (6), has to hold each period. Revenues from labor income taxes, τ_{lt}, and carbon taxes, τ_{Ft}, have to be at least as high as an exogenous funding requirement, Gov_t^{min}, plus spending on

[20]Feasibility means that the government is constrained by initial levels of knowledge and expertise, time endowments of workers and scientists, and production processes prescribed by the model.

research subsidies, $T_{Rt} = \tau_{sFt} w_{st} S_{Ft} + \tau_{sGt} w_{st} S_{Gt}$. Government consumption is modeled as a direct transfer to households.[21]

In the counterfactual version of the model, lump–sum taxation is feasible. Formally, this corresponds to relaxing the lower bound on revenues raised through distortionary instruments, so that the constraint (7) never binds. This counterfactual replicates the standard assumption in the directed technical change and climate–policy literature. Note that even if equipped with lump-sum taxes, the social planner allocation, defined in Appendix C, is infeasible due to learning-by-doing which additionally necessitates sector-specific technology subsidies.

3 Theory

This section establishes analytical results that guide the interpretation of the quantitative findings. The objective is to analyze how carbon taxation affects the allocation of research effort, and, in particular, how learning-by-doing shapes this relation. In settings with distortionary fiscal constraints, such indirect general-equilibrium effects are potentially valuable: if carbon taxes redirect research in a way the planner intents, the reliance on distortionary fiscal instruments to fund research subsidies reduces.

The following analysis centers on an *expertise effect*: when production shifts toward green energy, learning-by-doing raises green expertise relative to fossil expertise, thereby raising the returns to green innovation. Under plausible parameter values, this mechanism amplifies the impact of carbon taxation on the direction of research.

3.1 Equilibrium allocation of researchers and the expertise effect

In equilibrium, the demand for researchers equates marginal private costs and marginal private benefits. Assuming equal capital shares across sectors, $\alpha_G = \alpha_F = \alpha$, the ratio of

[21]The assumption ensures that government consumption generates household utility. If resources were wasted, instead, government consumption would affect households' labor supply decision.

green-to-fossil researchers can be written as[22]

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = \underbrace{1-\widehat{\tau}_{sFGt}}_{\text{research subsidies}} \underbrace{\left(\frac{p_{Gt}}{p_{Ft}}\right)^{\frac{1}{1-\alpha}}}_{\text{prices}} \underbrace{\frac{L_{Gt}}{L_{Ft}}}_{\text{market size}} \underbrace{\left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)}}_{\text{knowledge spillovers}} \underbrace{\left(\frac{A_{Gt}/K_{Gt}}{A_{Ft}/K_{Ft}}\right)}_{\text{marginal productivity}} \tag{8}$$

where $1-\widehat{\tau}_{sFGt} = \frac{1-\tau_{sFt}}{1-\tau_{sGt}}$ is the relative wedge between fossil and green subsidies. This expression highlights the key determinants of the direction of innovation: subsidies, prices, market size, knowledge spillovers, and the marginal productivity gains from new technologies. The latter effect is missing in papers assuming a one-to-one linkage between research and productivity (e.g. Acemoglu et al., 2012; Fried, 2018).

Substituting equilibrium conditions and the definition of productivity, eq. (4), yields the following equilibrium relation:

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = (1+\widehat{\tau}_{Ft})^{\varepsilon_e}(1-\widehat{\tau}_{sFGt}) \left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)} \underbrace{\left(\frac{Q_{Gt}}{Q_{Ft}}\right)^{\iota_L(1-\alpha)(\varepsilon_e-1)}}_{expertise\,effect} \left(\frac{K_{Gt}}{K_{Ft}}\right)^{\iota_K(1-\alpha)(\varepsilon_e-1)-1} \tag{9}$$

where $\widehat{\tau}_{Ft} = \frac{\tau_{Ft}}{p_{Ft}}$ is the carbon tax expressed as a fraction of the producer price of fossil energy.

The equation highlights that learning-by-doing gives rise to an expertise effect that directs innovative activity in general equilibrium operating through a market size effect: consider a shift in energy production toward green energy. Then, green expertise, Q_G, accumulates faster than fossil expertise, Q_F. Higher relative expertise raises labor productivity with elasticity ι_L, which increases the labor share of green production and lowers the relative price between green and fossil energy. When green and fossil goods are gross substitutes, $\varepsilon_e > 1$, as suggested by empirical evidence (Papageorgiou et al., 2017), the increase in output outweighs the decrease in prices and labor supply rises in total by an elasticity of $\iota_L(1-\alpha)(\varepsilon_e-1)$, feeding back into research incentives. Intuitively, research

[22] For the derivation of this and the following results see Appendix B. This section restricts attention to the special case with symmetric factor shares to simplify expressions and isolate the directive role of expertise. Derivations assume asymmetric factor shares. Allowing for asymmetric factor shares does not alter the direction of the mechanisms characterized here.

is directed toward the sector with more experienced workers. Proposition 1 summarizes this result.

Proposition 1 (**Expertise effect**). *Research is directed to the sector with more experienced workers if returns to research are decreasing, $\eta < 1$, the elasticity of productivity to expertise is positive, $\iota_L > 0$, the capital share is below unity, $\alpha < 1$, and energy inputs are gross substitutes, $\varepsilon_e > 1$. The elasticity of the ratio of green-to-fossil research to the ratio of green-to-fossil expertise is given by $\frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta}$.*

Proof. The proof follows directly from eq. (9). □

3.2 Expertise shapes the effectiveness of the carbon tax

Carbon taxation affects the allocation of researchers directly through a demand channel: by raising the relative price of fossil goods, demand shifts toward green goods. When green and fossil energy are gross substitutes ($\varepsilon_e > 1$), the carbon tax reallocates demand with elasticity above unity, increasing the market size of the green sector and thus the return to green innovation. The direct effect unambiguously spurs research in the green sector since it purely operates through shifts in demand.

The strength of this direct effect depends on the initial ratio of expertise. When green expertise is already relatively high, then a marginal increase in the carbon tax, ceteris paribus, implies a more pronounced shift to green research. Proposition 2 formalizes this observation.

Proposition 2 (**Amplification of direct carbon-tax effect via green expertise**). *A higher level of relative expertise in the green sector a-priori increases the marginal effect of the carbon tax on the ratio of green-to-fossil research, ceteris paribus, if $\frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta} > 0$.*

Proof. See Subsection B.1. □

A more effective carbon tax could, in principle, ease the burden on fiscal instruments by directing research to the green sector lowering the need for costly subsidies. Whether this is ultimately beneficial, however, depends on how well the direction in which the carbon tax pushes research aligns with the direction the planner prefers.

3.3 Indirect effect of carbon taxes through learning-by-doing

Beyond the direct effect described above, carbon taxation also influences the allocation of researchers through an indirect channel via learning-by-doing. To characterize this channel, I apply the Implicit Function Theorem to the log-transformation of eq. (9); see Subsection B.1. This yields the responsiveness of the research allocation to the carbon tax. Let $s_t = \log(S_{Gt}/S_{Ft}), k_t = \log(K_{Gt}/K_{Ft}), q_t = \log(Q_{Gt}/Q_{Ft}), \beta_Q = \frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta}$, and $\beta_K = \frac{\iota_K(1-\alpha)(\varepsilon_e-1)-1}{1-\eta}$, then

$$\frac{ds_t}{d\widehat{\tau}_{F,t}} = \frac{\frac{\varepsilon_e}{(1-\eta)(1+\widehat{\tau}_{F,t})} + \beta_Q \frac{\partial q_t}{\partial \widehat{\tau}_{F,t}}}{1 - \beta_K \frac{\partial k_t}{\partial s_t} - \beta_Q \frac{\partial q_t}{\partial s_t}}. \tag{10}$$

With learning-by-doing, a higher carbon tax raises green-to-fossil output and thus increases relative green expertise, implying $\partial q_t/\partial\widehat{\tau}_{F,t} > 0$. This term enters the numerator of eq. (10) and reinforces the effect of the carbon tax whenever $\beta_Q > 0$. Expertise also responds to the research allocation ($\partial q_t/\partial s_t > 0$), generating the feedback term in the denominator and further amplifying the adjustment of s_t. If expertise is exogenous, both derivatives are zero and the indirect effect vanishes.

Proposition 3 (**Indirect effect of carbon taxation through learning-by-doing**). *Carbon taxation indirectly affects the allocation of researchers through its effect on expertise accumulation. The marginal effect of the carbon tax on the research ratio is*

$$\frac{ds_t}{d\widehat{\tau}_{F,t}} = \frac{\frac{\varepsilon_e}{(1-\eta)(1+\widehat{\tau}_{F,t})} + \beta_Q \frac{\partial q_t}{\partial \widehat{\tau}_{F,t}}}{1 - \beta_K \frac{\partial k_t}{\partial s_t} - \beta_Q \frac{\partial q_t}{\partial s_t}}.$$

The indirect effect strengthens the overall impact of carbon taxation whenever research returns are decreasing, $\eta < 1$, expertise raises productivity, $\iota_L > 0$, the capital share is below unity, $\alpha < 1$, and green and fossil inputs are gross substitutes, $\varepsilon_e > 1$.

Proof. See Subsection B.1. □

Taken together, Proposition 2 and Proposition 3 show that learning-by-doing shapes how carbon taxation influences research incentives through two complementary channels.

The amplification of the direct effect implies that carbon taxes become more powerful in shifting private research toward the green sector when green expertise is already high. The indirect effect reflects that a carbon-tax-induced shift in production raises green expertise and thereby increases the marginal private return to green research. These channels illustrate how learning-by-doing can reduce the requirement to fund research subsidies by strengthening the research shifts induced by carbon taxation.

The existence of an innovation–learning dilemma, however, is a quantitative question that depends on how well the implied transition of researchers coincides with the target transition of research. Fiscal distortions become relevant when the planner seeks to disentangle the transition of production and expertise on the one hand and research on the other hand. Subsection 5.3 examines this question.

4 Calibration

This section describes the calibration of the model to U.S. data. I use a hybrid approach by assuming the economy was on a BGP up to and including 2010–2014 while allowing the base period 2015-2019 to be off the BGP. This flexibility enables capturing current productivity ratios in the data while anchoring structural parameters to long-run data.

In the main experiment, time-varying policy interventions to satisfy the emission target mean the economy is no longer on a BGP with constant ratios. I first discuss the data sources for the emissions target, stocks of technology and expertise, and emissions in Subsection 4.1. Subsection 4.2 outlines the parameter calibration procedure. Finally, Subsection 4.3 discusses the key calibrated parameters and compares non-target model outcomes to the data.

4.1 Data

Emissions target I focus on CO_2 emissions and omit other greenhouse gases as they are the largest contributor to warming (IPCC, 2023, p. 29). For the global CO_2 emissions pathway, I use the net emission budget consistent with a 2°C warming from the IPCC's sixth assessment report (Van der Wijst et al., 2023, Figure SPM.5). To derive a U.S.-specific target path, I adopt an *equal-per-capita* burden-sharing approach which is one of the principles discussed in Robiou du Pont et al. (2017). Specifically, I allocate the global carbon budget to the U.S. in proportion to its projected population share based on United Nations population forecasts (United Nations, 2022). This yields a declining pathway for U.S. net CO_2 emissions, $\{\Omega_t\}_{t\geq 2020}$, measured in gigatons per 5-year model period.

I find that meeting the target requires a 70.9% reduction of U.S. net emissions by 2030 (relative to 2019) and net-zero emissions starting in 2070 to limit warming to 2°C. For comparison, the Biden administration foresaw a 50-52% reduction by 2030 relative to 2005, which effectively corresponds to a 38% reduction relative to 2019.[23]

Knowledge stocks Calibration of the innovation side hinges on measuring the stocks of *green* and *fossil* energy knowledge—both their levels around the base period and their growth over time. The main challenge is classifying technological knowledge by energy type. To this end, I construct a novel time series of patents in each category using the universe of patents filed by U.S. applicants and granted by the USPTO sourced from the European Patent Office's (EPO) PATSTAT database. I filter technologies into green and fossil-related based on a taxonomy developed by the International Energy Agency and the EPO.[24] Fossil energy patents relate to the supply, exploration, processing, transport, and distribution of fossil fuels.[25] For green energy, I collect patents on low-carbon-energy supply technologies. The approach combines a two-step filtering procedure using, first, Cooperative Patent Classification (CPC) codes and, second, text-based filtering. This approach yields annual series for green and fossil patent applications from 1950 to

[23]Source: https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/22/, retrieved 14 September 2022.

[24]The classification tables are available in an EPO study on energy patents: see `https://link.epo.org/web/patents_and_the_energy_transition_study_en.pdf` for green technologies and `https://link.epo.org/web/patents_and_the_energy_transition_study_annex_en.pdf` for fossil-related technologies.

[25]These are technological advances which increase the output of fossil fuels from the same amount of inputs, hence, making fossil energy cheaper while emissions per unit of energy remain unchanged consistent with the model.

2014 for granted patents.[26] Knowledge stocks result as the cumulative patent counts accounting for depreciation. Subsection E.1 details the data collection and the derivation of knowledge stocks.

According to this data, the U.S. fossil sector's knowledge stock was about 25% larger than that of the green sector in the 2010-2014 period. The energy transition in this paper, thus, starts from an initial technological gap favoring fossil-fuel technologies. In contrast to previously found values in the literature, this measure is relatively low, reflecting the ongoing transition. Acemoglu et al. (2016) find a difference of 48%.[27] Fried (2018), who derives the initial distribution of knowledge from output data, finds a much higher knowledge advantage in the fossil sector of 250%.

Sectoral output and expertise Time series for green and fossil energy stem from (EIA, 2023, Table 1.1). I treat nuclear and renewables as the model's *green* energy and fossil fuels as *fossil* energy. The model's concept of expertise, Q_{Jt}, is not directly observed. I use cumulative sector-specific output as a proxy following the literature.[28] However, such measures depend on the unit of measurement of green and fossil energy production in the data, and there is no guarantee that these units align with the model variables, which are expressed as shares of base-period GDP. To derive a time series of expertise that is both consistent with the model and informed by historical sector-specific output growth, I proceed by combining level information from the model and growth information from the data. This approach yields synthetic time series of sectoral output consistent

[26] I focus on granted patents to ensure a certain quality of patents. The number of patents may not be a good proxy for *knowledge*, as patents can differ in their quality. An alternative measure used in the literature are citation-weighted patents which gives an idea about the stimulating force of an innovation. The more frequent a patent is cited, the more important the knowledge conveyed in this innovation. However, citation data most likely captures structural transformation processes and the green transition in particular. A fossil-related innovation, for instance, may see less citations not because it is of lower quality, but because all innovation happens in the green sector due to political intervention. This would understate the potential of fossil knowledge. Underestimating fossil-based knowledge, in turn, would lower the need for policy intervention to counter path-dependency of innovation. Using stock exchange information as used in Kogan et al. (2017) would also capture market expectations on policies and the greening of the economy, thus most likely understating knowledge advances in the fossil sector. The number of patents as a measure of knowledge relies on the assumption that the quality of patents within sectors is equal on average.

[27] This is the weighted average of knowledge stocks in clean and fossil sectors found in Acemoglu et al. (2016).

[28] See for instance Irwin and Klenow (1994).

with the model. They form the basis to construct expertise stocks as cumulative output in line with eq. (3). One advantage of this approach is that model-consistent time-series of expertise are available allowing to simulate the model from an arbitrary point back until the 1970s. For further details on the construction consult Subsection E.2.

Emissions data and sink capacity The Environmental Protection Agency's Inventory Greenhouse Gas Emissions and Sinks dataset (EPA, 2022) provides annual CO_2 emissions from energy, industrial processes and product use, and agriculture, as well as natural sinks such as forests. The dataset records total national emissions based on a territorial inventory meaning that emissions that physically occur in the U.S. are respected.[29] I use the data for calibrating both the sink capacity, δ, and the emissions-intensity of fossil production, ω.

The sink capacity follows directly from the data. Note that in the model, I treat the fossil energy sector as the sole source of CO_2. Therefore, I set the sink capacity to match the negative emissions of natural sinks net of emissions from industry and agriculture in the base period from 2015 to 2019. I, thus, implicitly fix emissions from other sectors assuming that all emissions reduction has to arise in the energy sector consistent with the model.[30] Using the EPA figures, the sink capacity is approximately $\delta = 3.19$ $GtCO_2$ per 5-year period, which I hold constant over time. This means that about 0.64 $GtCO_2$ per year are offset by natural processes.

4.2 Model parameters and calibration strategy

To calibrate the model parameters, I proceed in three steps. First, I assign certain parameters standard values from the literature or choose normalizations that set units. Second, I calibrate the *non-growth* side of the model (preferences, production, and baseline output levels) by matching key steady-state ratios and base-period observations. Third, I calibrate the *growth* side (innovation and learning parameters) using a target grid-search procedure, ensuring the model reproduces empirical patterns in knowledge stocks.

[29] For an overview see `https://www.epa.gov/ghgemissions/carbon-dioxide-emissions`.

[30] More formally: $-\delta =$ Natural sinks - Emissions from Industry and Agriculture.

Externally set parameters For household preferences, I set the 5-year discount factor to $\beta = 0.985^5 = 0.93$ (Barrage, 2020). The Frisch elasticity is set to 0.75 (Chetty et al., 2011), implying $\sigma = 1.33$. The disutility of labor χ is calibrated such that average hours worked equal 34 percent of the daily time endowment (14.5 hours), based on OECD data (OECD, 2021). I normalize the time endowment, $\bar{H}$, and the supply of scientists, S, to one.

On the production side, I set the elasticity of substitution between energy and non-energy inputs to $\varepsilon_y = 0.05$, reflecting that energy is an essential input with limited substitutability (Hassler et al., 2016). For substitution between green and fossil energy, I use $\varepsilon_e = 1.8$, following macroeconomic estimates in Papageorgiou et al. (2017). This implies that green energy can substitute for fossil energy to a limited extent. Capital shares are calibrated to match labor compensation data from BEA input-output tables for 2015–2019. I classify NAICS sectors 21 and 324 as fossil energy (resource extraction and refining), yielding $\alpha_F = 0.75$. For the non-energy sector, I use $\alpha_N = 0.36$, consistent with aggregate labor income shares. The exogenous growth rate of non-energy productivity is set to $\gamma_N = 0.09$, corresponding to 9 percent output growth per 5 years (OECD, 2022).

The initial carbon tax is set to $\tau_{F2015-19} = 0$, reflecting the absence of a federal carbon price in the U.S. in 2015. I also set subsidies on fossil research, $\tau_{sF2015-19} = 0$, and calibrate subsidies on green R&D as only the difference can be determined by model equations. The knowledge depreciation rate is set to $\delta_K = 0.55$ per period, or 15 percent annually (Noailly and Smeets, 2015). Finally, I match the elasticities of labor-augmenting productivity to knowledge and workers' expertise, $\iota_K = 0.21$ and $\iota_L = 0.50$ to the numbers found in Furman et al. (2002) and Irwin and Klenow (1994) as described in Subsection E.3.

Non-growth parameters Next, I turn to calibrating the non-research side of the model using, first, a BGP-matching and, second, a base-period approach. I start with $\{\alpha_G, \chi, \delta_y\}$ leveraging BGP conditions and long-run averages of key ratios. I target the average ratios of fossil-to-green knowledge K_F/K_G and of expertise, Q_F/Q_G, which jointly determine the implied ratio of fossil-to-green productivity, A_F/A_G, on the BGP. Finally, I discipline the model by matching the fossil-to-green energy ratio, F/G on the BGP from

(EIA, 2023, Table 1.1) as described above. I proxy the BGP-ratios as averages taken over the period 1990-2010.

I employ the average expenditure share on energy relative to GDP equal to 6% from the U.S. Energy Information Administration (EIA, 2023, Table 1.7) which informs the weight on energy in production.[31] The resulting weight on energy is $\delta_y = 0.33$.[32] The disutility from labor, χ, is calibrated so that average hours worked equal 34% of the daily available time (14.5 hours), based on OECD data (OECD, 2021). I derive an estimate of the labor share in the green energy sector from the green job tables of the Bureau of Labor Statistics (BLS).[33] The share of green energy employment to total employment equals 0.48% rationalized by a green capital share of $\alpha_G = 0.88$. By comparison, Fried (2018) uses a slightly higher value of 0.91.

Equipped with the parameter values that result from the BGP matching, I use base-year information to calibrate the labor tax and the emissions intensity of fossil energy $\{\tau_{l2015-19}, \omega\}$. To match $\tau_{l2015-19}$, I use observed government expenditures minus new debt relative to GDP in 2015-2019,[34] rationalizing $\tau_{l2015-19} = 0.26$; a value consistent with the literature (Barrage, 2020). To find the parameter relating CO_2 emissions[35] from energy use in the model, ω, I match gross emissions in the base period according to $Emissions = \omega F_{2015-19}$. The resulting emissions intensity of fossil-fuel production is $\omega = 191$. Base-period productivity levels $\{A_{F2015-19}, A_{N2015-19}, A_{G2015-19}\}$ follow from normalizing total output in 2015-19, $Y_{2015-19}$, to unity, taking it as the numeraire: $p_y = 1$ in all periods, and the base-period ratio of fossil-to-green energy (EIA, 2023, Table 1.1) as explained in the data section. This concludes the calibration of the non-research block of the model.

[31] Since within a calibration block, parameters are jointly determined within the general equilibrium, the relationship between targets and parameters is, hence, presented heuristically.

[32] Note that δ_y qualifies as a measure of energy efficiency in the economy.

[33] Retrieved from https://www.bls.gov/green/home.htm, 06 September 2023. The data only contains information for the late 2000s.

[34] The data comes from the Federal reserve Bank of St. Louis; https://fred.stlouisfed.org/series/W068RCQ027SBEA.

[35] The data comes from EPA (2022) and is described in Subsection 4.1.

Growth parameters The remaining parameters govern technological change: $\{\gamma, \eta, \phi, \nu, \tau_{sG2015-19}, \kappa, K_{G2010-14}, K_{F2010-14}, Q_{G2015-19}, Q_{F2015-19}\}$, which I pin down using a grid-search algorithm. The procedure can be summarized as follows.

I start by choosing candidate values for the returns to scale to research, η, and the initial green R&D subsidy, $\tau_{sG2015-19}$. For each candidate pair $\{\eta, \tau_{sG2015-19}\}$, I require parameter values of knowledge spillovers, ϕ, the productivity of researchers, γ, and the fishing-out exponent, ν, to be consistent with (i) the ratio of fossil to green knowledge, K_F/K_G, on the BGP, (ii) the output ratio, F/G, on the BGP, and (iii) the relative growth rate of knowledge in the base period.

As a second step, I can calculate the allocation of researchers and base-period knowledge stocks using already calibrated parameters and the initial knowledge stocks in the 2010-2014 period, $\{K_{F2010-14}, K_{G2010-14}\}$, which I take directly from the data. I run a constrained minimization routine to find the allocation. The constraint is to choose a productivity scaling factor κ, which converts units of K_J and Q_J into total factor productivity, so that the model's implied expertise levels $Q_{F2015-19}$ and $Q_{G2015-19}$ are consistent with the initial output split while close to the levels derived from the data.

Intuitively, the model determines sectoral output levels and knowledge stocks in the base period. Fixing sectoral expertise stocks to the data would yield a system of two equations and one unknown that most likely has no solution. Therefore, I allow more flexibility and let the minimization routine determine κ so that the distance of model-implied and data-implied levels of expertise is minimized. As a result, $\{Q_{F2010-14}, Q_{G2010-14}\}$ are model implied, while $\{K_{F2010-14}, K_{G2010-14}\}$ are exactly matched to the data.

After solving for $\{\phi, \nu, \gamma, \kappa, Q_{F2010-14}, Q_{G2010-14}\}$ given a candidate $\{\eta, \tau_{sG2015-19}\}$, I simulate the model forward from 1970 until 2010 using initial stocks of knowledge and expertise from the data. I then select the set of parameters that minimizes the distance of the model to the observed evolution of green and fossil knowledge stocks.

4.3 Discussion

Before turning to evaluate the model against empirics, I discuss some important parameter values and how they relate to the literature. All resulting parameter values are depicted in Table 1.

Comparison to the literature The resulting relative importance of cross-sectoral knowledge spillovers is $\phi = 0.44$. Aghion et al. (2016) estimate for the U.S. automotive industry that clean innovation within a firm is comparably more important for clean patent growths than dirty knowledge. Matching the relative importance of within- to cross-sectoral spillovers from their estimation yields $\phi^{AAHK} = 0.3124$.[36] Since they focus on the automotive industry and micro-level estimates, they do not include spillovers across firms. Since my model accounts for fishing-out effects, the relevant absolute strength of cross-sectoral knowledge stock spillovers is given by $\phi(1-\nu) = 0.37$ which is well in the range of values in the literature.[37] Hart (2019) calibrates a value equivalent to $\phi(1-\nu) = 0.1$ and Fried (2018) sets $\phi(1-\nu) = 0.5$ based on theoretic considerations.

For the stepping-on-toes effect, I find a value of $\eta = 0.41$. The value below unity can be explained by the probability of duplicating ideas that rises the more scientists work on the same research process. Again, the value lies in the range of estimates used in the literature. Acemoglu et al. (2016) find a similar value of $\eta = 0.37$ in a first-difference estimation based on micro-level data on the energy sector. Fried (2018) estimates $\eta = 0.79$. The higher value implies that a less equal allocation of scientists is more productive than in the calibration presented here. Hart (2019), in contrast, finds a value of $\eta = 0.19$. Finally, fishing-out effects are non-negligible as suggested by the literature (Bloom et al., 2020; Kruse-Andersen, 2023) but small with a value of $\nu = 0.16$.

Model fit Figure 1 contrasts model outcomes (blue solid) to the data (orange dashed) for non-targeted variables. Figure 1a depicts gross emissions. The model captures dynamics of emissions fairly well until the early 2000. Afterwards, policy interventions

[36]They estimate an elasticity of new clean innovation to past clean innovation of 0.306 compared to an elasticity of 0.139 with respect to past dirty innovation.

[37]Within-sectoral spillovers amount to $(1-\phi)(1-\nu) = 0.46$.

Table 1: Calibration

Parameter	Main Target	Value
Household		
Inverse Frisch elasticity σ	Chetty et al. (2011)	1.33
Disutility from labor χ	average hours worked per economic time endowment by worker: 0.34 (OECD, 2021)	9.66
Discount factor β	Barrage (2020)	0.93
Working time endowment $\bar{H}$	14.5 hours per day (Jones et al., 1993)	1.00
Scientists S	normalization	1.00
Research		
Returns to research η	evolution knowledge stocks	0.41
Knowledge spillovers ϕ	evolution knowledge stocks	0.44
Scientists' productivity γ	evolution knowledge stocks	1.12
Fishing-out effect ν	evolution knowledge stocks	0.16
Initial knowledge stock ($K_{F2010-14}$, $K_{G2010-14}$)	knowledge stock in 2010-2014	(1.25, 1.00)
Initial expertise ($Q_{F2010-14}, Q_{G2010-14}$)	matching knowledge stock and output	(0.14, 0.07)
Elasticity of productivity (ι_L, ι_K)	(Irwin and Klenow (1994), Furman et al. (2002))	(0.50, 0.21)
Depreciation knowledge stock δ_K	Noailly and Smeets (2015)	0.55
Production		
Elasticities of substitution (ε_y, ε_e)	(Fried (2018), Papageorgiou et al. (2017))	(0.05, 1.50)
Weight on energy in output δ_y	expenditure share on energy (EIA, 2023)	0.33
Capital shares (α_F, α_G, α_N)	BLS and Green Jobs and Compensation of employees	(0.75, 0.88, 0.36)
Non-energy productivity growth γ_N	5-year GDP growth (OECD, 2022)	0.09
Government		
Policy instruments ($\tau_{F0}, \tau_{sF0}, \tau_{sG0}, \tau_{l0}$)	Barrage (2020) and knowledge stocks	(0, 0.64, 0, 0.26)
Emissions		
Carbon sinks δ in $GtCO_2$	EPA (2022)	3.19
Emissions per unit of fossil energy ω	EPA (2022)	190.85

or non-modeled features implied a structural break in the data leading to a downward sloping path. Absent policy adjustments, the model predicts a rise in gross emissions amounting to almost 7$GtCO_2$-equivalent in 2010, which exceeds observed emissions by 27%.

Figure 1b presents the comparison of the model and data-implied ratio of fossil-to-green output. The model replicates the steady decline in the ratio of fossil-to-green energy since the 1970s, though the decline is somewhat muted relative to the data. The non-modeled oil-price shocks of the 1970s and regulatory changes are candidate explanations.

Taken together, these figures provide reassurance that the model captures both the broader emission dynamics and the underlying trends in knowledge accumulation reasonably well.

Figure 1: Model Evaluation

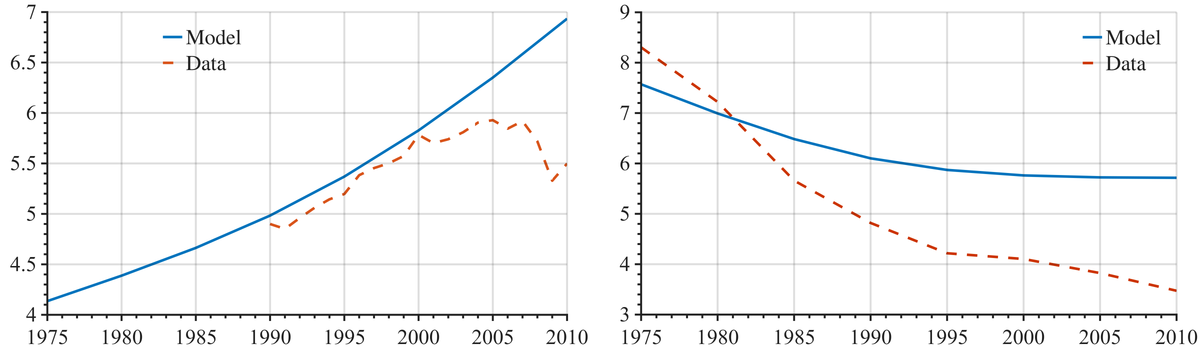

(a) Annual gross emissions **(b)** Fossil-to-green output

Notes: Panels 1a and 1b report annual gross emissions and fossil-to-green energy output ratio from the model and from the data. Both dynamics are non-targeted in the calibration. The model assumes a constant emission intensity of fossil energy. To account for structural changes in emission intensity per unit of fossil energy ω, I recalibrate this parameter so that the model matches emissions in the 1990-1994 period which is the first period for which I have emissions data. Importantly, this does not affect dynamics of emissions.

5 Results

This section presents the main quantitative results. I first describe the optimal policy mix in the baseline model and its implications for the composition of output and research along the transition path in Subsection 5.1. I then analyze how fiscal distortions and learning-by-doing shape the optimal fiscal mix in Subsection 5.2. Next, Subsection 5.3 examines the implications of fiscal distortions and learning-by-doing for the allocation of innovation. Finally, I assess the welfare implications in Subsection 5.4 and discuss sensitivity of the main results in Subsection 5.5.

5.1 The optimal policy mix

Figure 2 summarizes the main quantitative results. Figure 2a shows the carbon tax relative to the social cost of carbon (SCC), and Figure 2b reports the optimal carbon tax in levels. Figure 2c and Figure 2d depict the optimal labor income tax and spending on research subsidies to GDP both expressed as deviations from their optimal levels in an otherwise identical economy without emissions constraint. Figure 2e and Figure 2f display the implied allocation in levels, reflected by the green-to-fossil energy ratio and the fossil-to-green R&D ratio.

The optimal policy relies heavily on carbon pricing. As shown in Figure 2b, the carbon tax rises from an initial level of \$60 to about \$480 by 2070 (2022 prices). The carbon tax exceeds the SCC[38] by about 3.5% initially (Figure 2a). The wedge narrows to roughly 2.5% by 2030 and remains positive throughout the transition, despite the presence of fiscal distortions.

The labor income tax declines relative to its non-target optimum (Figure 2c). Initially, the reduction amounts to 4.3 percentage points and gradually falls to 2.2 percentage points by 2075. This pattern reflects the recycling of carbon tax revenues to reduce pre-existing labor tax distortions. Over time, the reduction becomes smaller as carbon tax revenues shrink whereas government spending requirements grow with output.

Expenditures on research subsidies relative to GDP are substantially reduced in the early phases of the transition. Relative to the non-target optimal level, R&D subsidy spending falls by 69% initially (Figure 2d). Only around 2040–2045 do research subsidies exceed their non-target level. In the long run they are about 90% higher than in the non-target optimum.

Under this policy mix, the emissions target is exactly met. The green-to-fossil energy ratio rises monotonically over time (Figure 2e), while the fossil-to-green R&D ratio declines smoothly from about 0.8 initially to roughly 0.65 in the long run (Figure 2f).

5.2 Mechanism: The optimal fiscal mix

This section clarifies how the optimal policy mix should be designed as part of fiscal policy and how learning-by-doing modifies second-best policies. The analysis decomposes the role of fiscal distortions into a component that operates when learning dynamics are taken as given and a component that arises once learning-by-doing itself becomes a policy objective. I proceed in three steps. First, I analyze a benchmark economy with lump-sum taxation and endogenous learning-by-doing. Second, I introduce distortionary taxation while holding sectoral expertise fixed at the path induced under lump-sum finance, so that

[38] I define the SCC as the shadow value of the emissions constraint in the Ramsey planner's problem divided by the household's marginal utility of consumption. The carbon tax coincides with the so-defined SCC in a model with exogenous expertise accumulation and lump-sum taxation.

Figure 2: Optimal policy mix

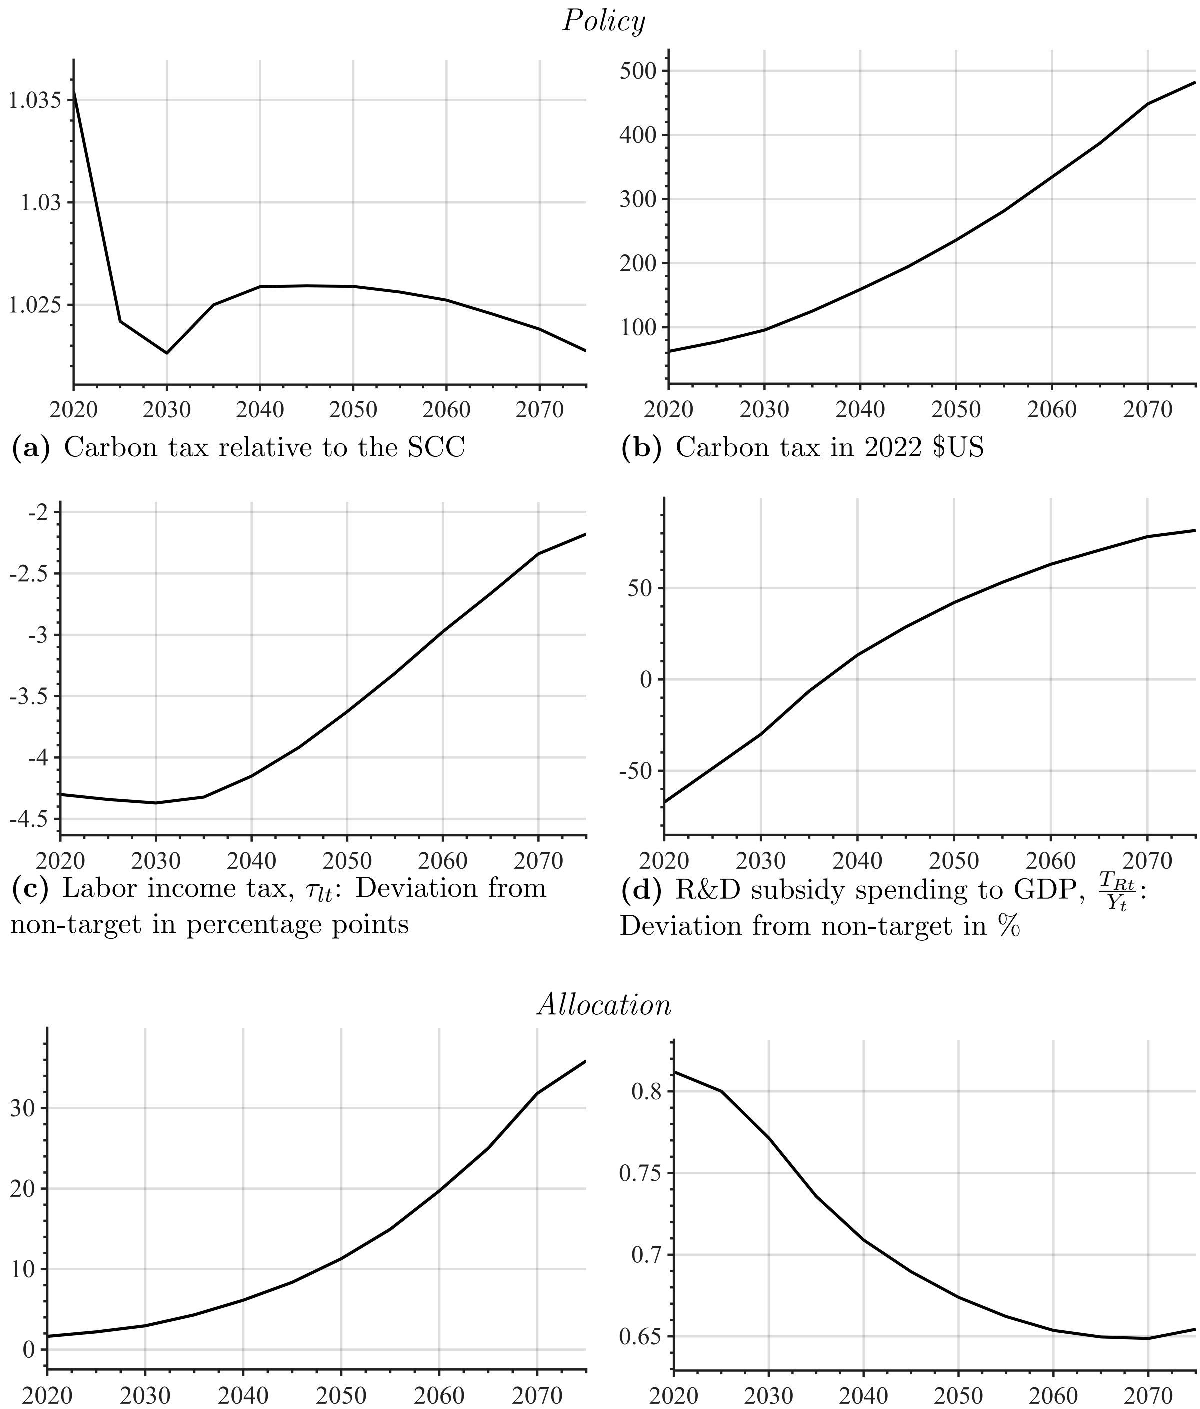

(a) Carbon tax relative to the SCC

(b) Carbon tax in 2022 \$US

(c) Labor income tax, τ_{lt}: Deviation from non-target in percentage points

(d) R&D subsidy spending to GDP, $\frac{T_{Rt}}{Y_t}$: Deviation from non-target in %

(e) Green-to-fossil energy ratio, $\frac{G_t}{F_t}$, in levels

(f) Fossil-to-green R&D ratio, $\frac{S_{Ft}}{S_{Gt}}$, in levels

Notes: The figure reports the optimal time paths for policy instruments and allocations under the baseline calibration. Panels 2a–2b show the carbon tax relative to the SCC and in levels. Panels 2c–2d report the labor income tax and research subsidy spending to GDP both expressed as deviations from the optimal level in an otherwise identical economy without emissions target. Panels 2e–2f display the optimal ratio of green-to-fossil energy and of fossil-to-green research effort under the emissions target. The x-axis reports the first year of the five-year period to which each variable refers.

the planner faces fiscal distortions but expertise dynamics are taken as given. Third, I allow learning-by-doing to adjust endogenously under distortionary taxation, which yields the baseline model.

Figure 3 provides a visual summary. Figure 3a and Figure 3b depict the carbon tax relative to the SCC and in levels, respectively. Figure 3c reports the labor income tax expressed as deviations from its non-target optimum, and Figure 3d depicts the same instrument in levels. Figure 3e and Figure 3f report the implied green-to-fossil energy ratio and R&D subsidy spending in levels. Each panel displays three regimes: the lump-sum benchmark with endogenous learning (blue dashed), the distortionary fiscal regime with expertise held fixed at its benchmark path (orange dash-dotted), and the baseline with fiscal distortions and endogenous learning-by-doing (black solid).

Fiscal distortions: Mitigating fiscal distortions I begin by comparing a regime with fiscal distortions and fixed expertise to a non-distortionary benchmark. In the presence of distortionary taxation, the carbon tax required to meet the emissions target is lower in levels (Figure 3b) due to pre-existing tax distortions that reduce economic output in general. Furthermore, the carbon tax lies below the SCC (Figure 3a). This reflects the classic tax-interaction effect in second-best settings going back to Bovenberg and De Mooij (1994). The key insight is that carbon taxes tend to be more distortionary than broad-based labor income taxes because they not only shrink their own tax base but also erode the labor-tax base by lowering real wages and labor supply.[39] Hence, raising revenue through carbon taxation imposes additional efficiency costs compared to labor taxation, and a second-best planner limits the reliance on carbon pricing to avoid compounded distortions.

The negative adjustment in the labor tax (Figure 3c) financed by carbon tax revenues is in line with the weak double-dividend logic: carbon tax revenues are best recycled to reduce pre-existing labor tax distortions (Bovenberg and De Mooij, 1994; Goulder, 1995). Accordingly, the labor income tax falls relative to its non-target optimum along

[39] The effect on labor supply arises through general-equilibrium mechanisms: by shifting production to a less productive mix, carbon pricing lowers the marginal productivity of labor. The after-tax wage rate declines thereby depressing labor supply under common parameter values.

the transition (Figure 3c), with the reduction being largest early on when carbon tax revenues are highest.

As a result, pre-existing tax distortions alter the scale and composition at which the emissions target is implemented. Under distortionary fiscal policy, the target is achieved at a smaller output scale and a higher fossil share than under lump-sum finance, reflected in a lower green-to-fossil energy ratio (Figure 3e). This point relates to the perspective emphasized by Metcalf (2003): second-best carbon taxes can be lower without necessarily implying weaker environmental outcomes because the broader fiscal environment affects the equilibrium scale and composition of production. In Subsection 5.3, I show that this change in the composition of energy production affects the allocation of research the planner aims to implement.

Learning-by-doing: Accepting a more distortionary fiscal mix Next, I discuss how learning-by-doing modifies second-best policies by comparing the distortionary baseline (black solid) to the distortionary regime under fixed expertise (orange dash-dotted). Allowing expertise to adjust endogenously raises the optimal carbon tax relative to the fixed-expertise case and reverses the wedge to the SCC (Figure 3a). In the absence of targeted instruments to directly subsidize learning, carbon pricing becomes a tool to foster learning in the green sector.

This result is notable because the public finance literature emphasizes that learning-by-doing can increase the distortionary cost of fiscal instruments by affecting not only current labor supply but also future productivity (Peterman, 2016). Nonetheless, the planner optimally relies more heavily on carbon pricing, and less on labor taxation, despite the associated increase in labor market distortions. The reason is that, during the transition, the relevant margin is not solely aggregate output, but the direction of expertise accumulation, which shapes future comparative advantage in green production.

Figure 3: Fiscal distortions and learning-by-doing

Policy

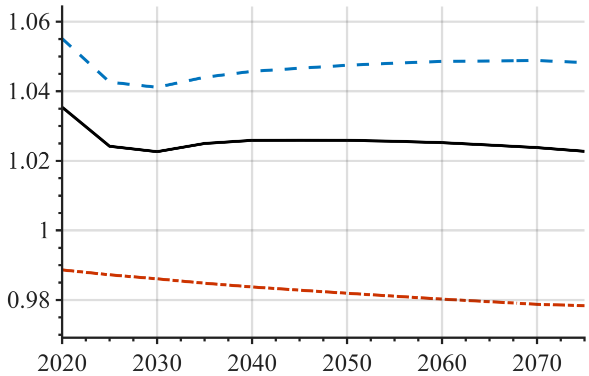

(a) Carbon tax relative to the SCC

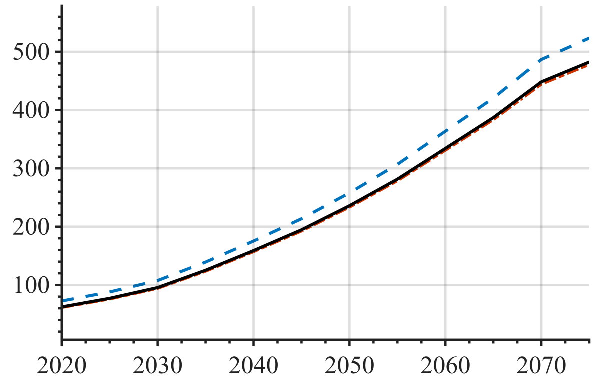

(b) Carbon tax in 2022 \$US

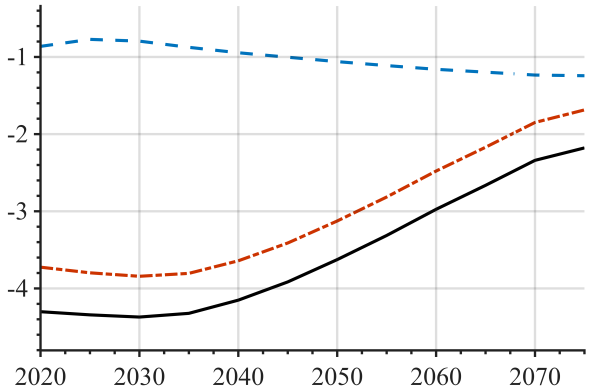

(c) Labor income tax, τ_{lt}: Deviation from non-target in percentage points

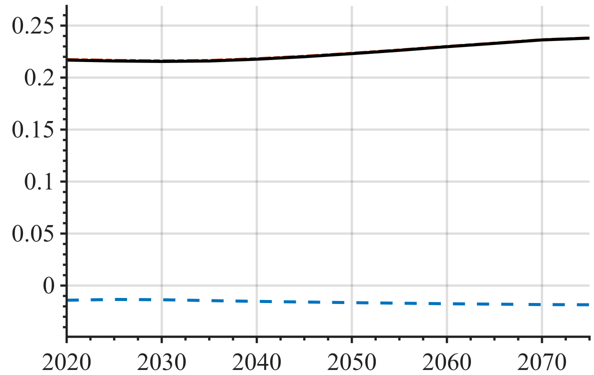

(d) Labor income tax, τ_{lt}, in levels

Allocation

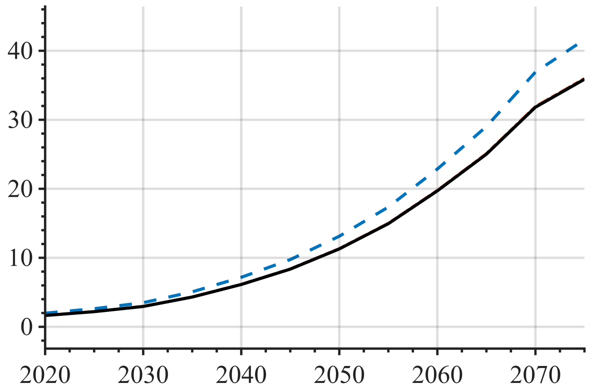

(e) Green-to-fossil energy ratio, $\frac{G_t}{F_t}$, in levels

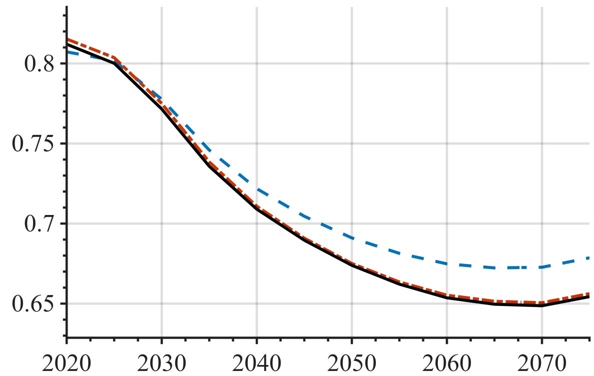

(f) Fossil-to-green R&D ratio, $\frac{S_{Ft}}{S_{Gt}}$, in levels

Notes: The figure compares main policy and allocation variables across three regimes: (i) a lump-sum benchmark with endogenous learning-by-doing (blue dashed graph), (ii) a regime with fiscal distortions and expertise fixed at the path that is implemented under lump-sum taxation (orange dashed-dotted graph), and (iii) the baseline with fiscal distortions and endogenous learning-by-doing (black solid graph). Panels 3a and 3b depict the carbon tax relative to the SCC and in levels. Panel 3c reports the labor income tax expressed as deviations from its non-target optimum, and Panel 3d depicts the labor income tax in levels. Panels 3e and 3f report the green-to-fossil energy ratio and the fossil-to-green ratio of scientists arising under the optimal emissions target implementation. The x-axis reports the first year of the five-year period to which each variable refers.

The results under the non-distortionary fiscal regime confirm this interpretation (blue dashed). Here, the carbon tax is between 4 to 5.5% higher than the SCC. In addition, the planner accommodates the excessive carbon tax with a subsidy on labor. The labor income tax is reduced by one percentage point, on average, relative to the non-target optimum (Figure 3c) and is negative in levels (Figure 3d).[40] In the distortionary regime, such a labor subsidy is infeasible. Nevertheless, the planner accepts the additional distortions induced by higher carbon taxation.

Dynamic importance of green expertise and innovation policy Turning to the dynamics of the wedge between carbon taxation and the SCC reveals that policies promoting green expertise are quantitatively most important early in the transition, indicating a high value of building expertise at initial stages. By contrast, research subsidies gain importance over time. These patterns point to a shift in policy emphasis: from instruments that primarily steer production toward green technologies to those that increasingly shape the direction of research.

While this section has focused on the optimal mix of carbon and labor taxation, fiscal constraints also affect the use of research subsidies and, in turn, the allocation of innovation, which I analyze next.

5.3 Mechanism: Misallocation of research

I now examine how distortionary fiscal constraints distort the allocation of innovation. The purpose of this section is to show how learning-by-doing alters the importance of fiscal constraints for the allocation of research. In particular, I show that the stronger carbon taxation to foster green expertise entails too fast a transition of innovative activity, motivating a more decisive use of research subsidies. At the same time, fiscal policies become more distortionary when learning-by-doing is endogenous and the planner accepts a more severe misallocation of research.

[40] Comparing these results to the social planner allocation shows that even without explicit instruments to target sectoral learning-by-doing the combination of an above-SCC carbon tax and a labor subsidy closely approximates the efficient outcome where learning spillovers are fully internalized.

Measuring misallocation of innovation To gauge the extent of misallocation of innovative activity, I first define a suitable benchmark. A naive comparison between the model with fiscal distortions and the non-distortionary setting would provide a misleading measure. The reason is that distortionary finance affects not only the feasibility of implementing a given research allocation, but also the allocation the planner would like to implement in the first place. As established previously, under fiscal distortions the emissions target is implemented at a higher share of fossil energy and a lower overall output level, most likely raising the social value of fossil research.

To disentangle the impact of changing target allocations and changing feasibility of innovation policy, I conduct the following experiment. I introduce a regime with lump-sum finance for R&D subsidies: the planner can finance green and fossil research subsidies using lump-sum revenue, while all other fiscal instruments remain distortionary. This setup preserves the broader fiscal environment and the resulting scale and composition of output—given the small size of research subsidies—while releasing the planner from rationing subsidies. I refer to the allocation implemented in this regime as the *target* allocation of research. Misallocation arising from fiscal policy is given by the difference between the fully distortionary regime and the target allocation.

The (target) allocation of researchers under fiscal distortions The experiment is informative on the role of fiscal constraints for research subsidy expenditures. Figure 4b shows R&D subsidy expenditures to GDP in levels under fiscal distortions (black solid), in the non-distortionary benchmark (blue dashed), and under lump-sum–financed subsidies (orange dash-dotted). Throughout the transition the planner lowers research subsidy expenditures relative to the non-distortionary benchmark due to a lower requirement; compare the non-distortionary case to the lump-sum–financed subsidy benchmark. In addition, the planner curtails subsidies to forgo fiscal pressures; compare the orange dash-dotted to the black solid graph. Consequently, the planner accepts some misallocation of innovation.

Figure 4a displays the allocation of researchers as the ratio of fossil-to-green scientists across regimes. The target allocation under fiscal distortions (orange dash-dotted) features a higher fossil R&D share than in the non-distortionary setting (blue dashed), reflecting the higher fossil share in energy production. Under fiscal distortions, the ratio increases by roughly 3 percentage points throughout the transition.

With distortionary fiscal policy, the planner optimally curtails subsidy expenditures reducing the fossil share beyond its target (black solid). Figure 4c quantifies the resulting misallocation. The black solid line shows the percentage deviation of the fossil-to-green research ratio under fiscal distortions from the corresponding ratio with lump-sum–financed subsidies. Misallocation ranges between 3.3% and 4.2% during the transition.

Complementarity of carbon taxes and research subsidies The misallocation of research arises because private and social benefits diverge as machine producers do not internalize the dynamic spillovers of knowledge. In particular, cross-sectoral spillovers make it socially commendable to maintain fossil research during the transition—not for its direct technological contribution, but because it enriches the knowledge base that supports green innovation.[41] The maintained social value of fossil research conflicts with the urgency to shift production toward green energy sources. As highlighted in the theory section, market forces direct research toward green technologies in response to higher carbon taxes.

Research subsidies help break this link. They optimally complement carbon taxes by slowing an excessively rapid reallocation of research toward green technologies. In a second-best setting, the carbon tax required to satisfy the emissions target is lower, which already dampens the shift of innovation toward the green sector. This observation helps explain the reduced reliance on research subsidies discussed above.

[41] Barbieri et al. (2023) provide empirical evidence for positive spillovers from fossil to green innovation.

Figure 4: Misallocation of innovation under fiscal distortions

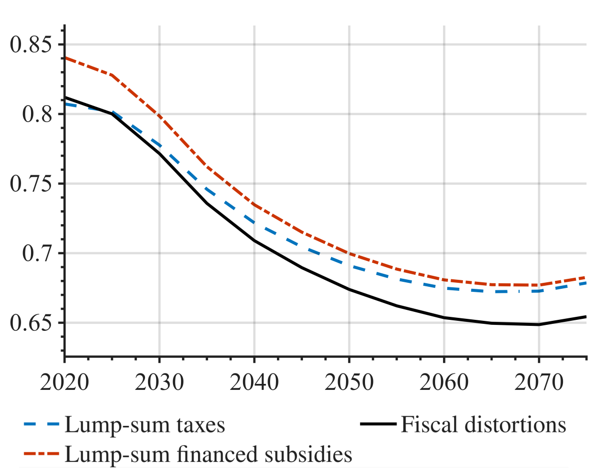

(a) Fossil-to-green R&D ratio, $\frac{S_{Ft}}{S_{Gt}}$

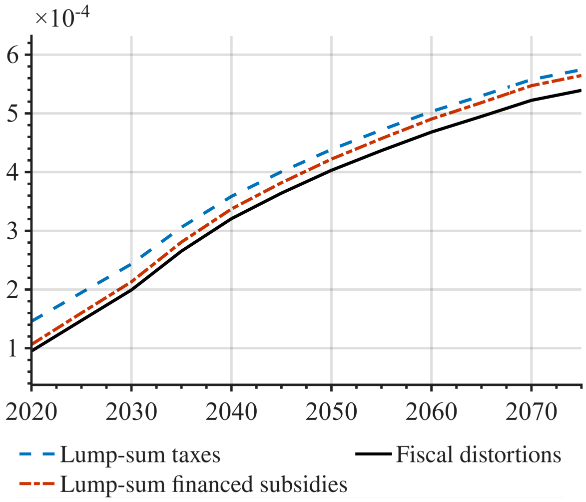

(b) R&D subsidy spending to GDP, $\frac{T_{Rt}}{Y_t}$

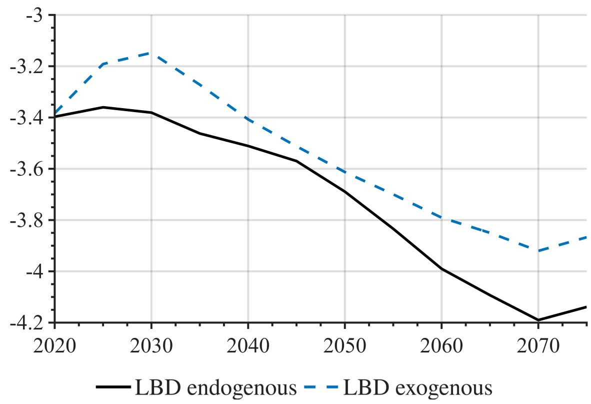

(c) Misallocation: Deviation from allocation under lump-sum–financed subsidies in %

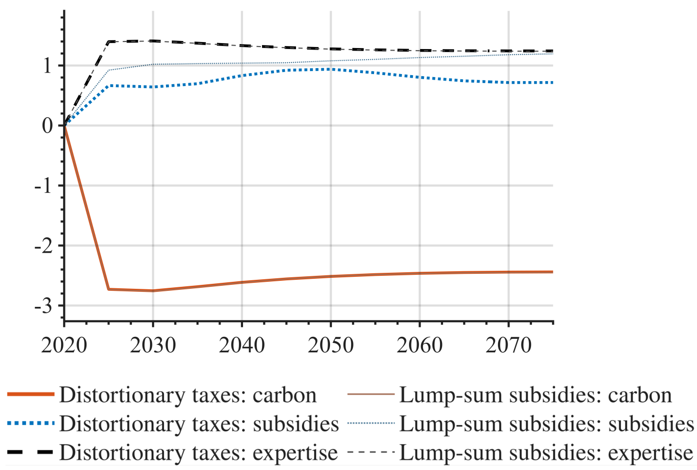

(d) Decomposition: Effect of endogenizing learning-by-doing on fossil-to-green R&D ratio

Notes: Panel 4a shows the ratio of fossil-to-green scientists under different fiscal regimes in the baseline model. For the same regimes, Panel 4b reports R&D subsidy expenditures relative to GDP in levels. Panel 4c quantifies the misallocation of innovation induced by fiscal distortions, measured as the percentage deviation of the fossil-to-green R&D ratio in the fully distortionary regime from a benchmark with lump-sum financed research subsidies for the model with and without endogenous learning-by-doing. Finally, Panel 4d decomposes the difference in the logged ratio of fossil-to-green research allocation between the model with endogenous and exogenous learning-by-doing into ceteris paribus contributions from carbon taxes, research subsidies, and expertise. Reported percentages are computed as $(\exp(x) - 1) \times 100$ and measure changes in the fossil-to-green research ratio relative to the exogenous-learning benchmark. The x-axis reports the first year of the five-year period to which each variable refers.

Learning-by-doing and the misallocation of innovation The same implication of carbon taxation is crucial for understanding how learning-by-doing affects the misallocation of innovation. As shown in Subsection 5.2, endogenous learning-by-doing motivates the use of higher carbon taxes (Figure 3b). While this policy response supports the accumulation of green expertise, it intensifies the conflict with the objective to maintain fossil research for its spillover benefits.

Figure 4c compares the baseline misallocation to the same metric arising in the model where learning is exogenous. It reveals that learning-by-doing amplifies the misallocation of innovation in all periods. Misallocation rises by between 4.4% on average when learning-by-doing is endogenous.

Decomposing the effect of learning-by-doing on the allocation of research To understand the underlying mechanisms, I return to the general-equilibrium condition that determines the allocation of research effort between fossil and green technologies.[42] The decomposition serves to quantify how different forces—carbon taxation, research subsidies, and endogenous expertise—contribute to the observed amplification of misallocation under learning-by-doing.

Taking logs of the inverted condition yields an additive expression for the research ratio, $\log(\frac{S_{Ft}}{S_{Gt}})$. Figure 4d exploits this additive structure to quantify the ceteris paribus effect of (i) the carbon tax, $(1 + \frac{\tau_{Ft}}{p_{Ft}})$, (ii) relative research subsidies, $\frac{1-\tau_{sGt}}{1-\tau_{sFt}}$, and (iii) relative expertise, $\frac{Q_{Gt}}{Q_{Ft}}$ on the change in research allocation induced by endogenizing learning-by-doing, that is, their contribution to the difference $\log(\frac{S_{Ft}}{S_{Gt}})^{\text{endo LBD}} - \log(\frac{S_{Ft}}{S_{Gt}})^{\text{exo LBD}}$. Effects are given in percent relative to the fossil-to-green ratio of scientists under exogenous learning-by-doing. The thick lines refer to the distortionary fiscal regime.

First, higher carbon taxes required to sustain green learning shift research toward the green sector, reducing the fossil-to-green research ratio by between 2.5 and 2.8% (orange solid). Second, the expertise effect works in the opposite direction: because fiscal distortions imply a less green production mix, the share of fossil research rises by slightly

[42] The decomposition builds on the general-equilibrium condition governing the allocation of research effort. Eq. (9) provides a transparent analytical benchmark; in the quantitative analysis presented here, I use the exact general form as stated in Appendix B, evaluated at the calibrated parameter values.

more than 1% (black dashed).[43] Third, research subsidies adjust to counteract the higher carbon tax, but at less than 1% (blue dotted).

The thin lines in Figure 4d show the same decomposition under lump-sum–financed subsidies, corresponding to changes in the target allocation of research. In this case, subsidies would respond more strongly to offset the carbon-tax-induced shift toward green research. Thus, when learning-by-doing is endogenous, the planner would like to use more research subsidies to decouple the transition of innovation from the transition of expertise, but fiscal constraints prevent their sufficient use.

While the effect of fiscal policy on the (mis)allocation of innovation is persistent and policy-relevant, it is not the primary driver of the welfare costs as quantified next.

5.4 Welfare costs of fiscal distortions

This section reinterprets the mechanisms discussed above from a welfare perspective. I quantify the welfare cost of distortionary fiscal policy and how it is shaped by learning-by-doing using consumption-equivalent variation (CEV).[44] I compute welfare costs under both endogenous and exogenous expertise accumulation and organize the analysis around two dimensions. First, I compare welfare costs of distortionary fiscal policy in the presence and absence of an emissions target to contrast the implications of learning-by-doing across scenarios. Second, I compare environments with and without lump-sum–financed research subsidies to isolate the distortions arising from the need to direct research under fiscal constraints.

The results are summarized in Table 2. When an emissions target is imposed, the welfare loss from fiscal distortions amounts to 1.3% of per-period consumption under exogenous learning-by-doing and rises to 1.5% when learning-by-doing is endogenous,

[43] By Proposition 2, this change in expertise mitigates the direct effect of carbon taxation on the allocation of research. The direction of this effect depends on the experimental baseline. Starting from a lower green-to-fossil expertise ratio, the same expertise channel would instead amplify the effect of carbon taxation on research allocation. Similarly, in the absence of cross-sectoral knowledge spillovers, the planner targets an even greener research allocation. In that environment, a relatively low green-to-fossil expertise ratio complicates attaining this target. Subsection F.1 depicts the analysis for this case.

[44] The CEV is defined as the permanent proportional reduction in per-period consumption under lump-sum taxation that makes the representative household indifferent between the non-distortionary and distortionary fiscal regimes.

Table 2: Consumption-equivalent variation decomposition

	Exo. LBD	**Endo. LBD**	**Effect of LBD (%)**
Panel A: Emissions target			
With emissions target	1.300%	1.499%	15.30%
Without emissions target	1.710%	1.850%	8.00%
Panel B: Innovation financing channel			
Innovation financing distortions†	0.0054%	0.0059%	10.55%
No innovation financing distortions	1.295%	1.494%	15.32%

Notes: This table reports consumption-equivalent variation (CEV), defined as the permanent proportional reduction in per-period consumption under lump-sum taxation that makes the representative household indifferent between the non-distortionary and distortionary fiscal regimes. The *Effect of LBD* column reports the percentage change induced by endogenous learning-by-doing (*Endo. LBD*) relative to the exogenous learning-by-doing (*Exo. LBD*) benchmark. In the latter, the path of expertise is fixed at the implemented path under lump-sum taxation.
† This row isolates the welfare cost associated with distortionary funding of research subsidies.

corresponding to a 15% amplification. Without an emissions target, the corresponding amplification is more modest, at about 8%. These results indicate that the amplification of fiscal distortions becomes substantially more costly during the transition to net-zero emissions.

This pattern is consistent with a trade-off intrinsic to learning-by-doing during the transition to net zero emissions. On the one hand, labor supply becomes more valuable because it jointly determines current production and future productivity. On the other hand, steering the direction of learning toward green technologies relies on distortionary fiscal instruments—i.e., the carbon tax.

It turns out that the misallocation of innovation induced by endogenous learning-by-doing is not driving the amplification. Panel B in Table 2 shows that eliminating distortionary financing of research subsidies leaves the 15% amplification essentially unchanged. Although learning-by-doing increases the welfare cost of this misallocation by 11%, the effect is small in levels. The overall amplification is therefore driven primarily by the interaction between distorted labor supply and the accumulation of sector-specific expertise.

5.5 Robustness

To assess the stability of the main quantitative results, I conduct a series of robustness exercises. Each experiment varies a single parameter while holding the rest of the calibration fixed. The analysis focuses on parameters governing (i) the labor share in the green sector, α_g, (ii) the elasticity of labor-augmenting productivity to learning-by-doing, ι_L, (iii) the strength of cross-sectoral knowledge spillovers, ϕ, (iv) the returns to research, η, and (v) policy targets, including the level of government expenditures, Gov_t^{min}, and the emissions constraint. For each specification, I report the optimal carbon tax relative to the SCC and the misallocation of innovation. Results are summarized in Table 3 in Appendix F.

Carbon taxation Across all robustness specifications, learning-by-doing raises the optimal carbon tax above the SCC throughout the transition (Panel A of Table 3). In every case, the wedge is most pronounced in initial periods and declines over time, consistent with the incentive to front-load the accumulation of green expertise. Importantly, within each period the carbon-to-SCC ratio is tightly clustered around its baseline value. Abstracting from changes in the labor share of green production, the carbon tax remains within ± 2.2 percentage points of the baseline across all parameter variations.

The primary source of dispersion across robustness specifications is the labor share in the green sector, $1 - \alpha_G$. When green production is more labor intensive at $\alpha_G = 0.75$, the optimal carbon tax exceeds the SCC by up to 35% in the short run. This pattern is consistent with the heightened value of green expertise when production relies more strongly on experienced workers. The baseline calibration, with a low labor share in green production corresponding to $1 - \alpha_G = 0.12$, therefore, represents a conservative parameter choice. Quantitative effects under the baseline calibration lie at the lower end of the effect size which appear to increase monotonically with the labor share in green production. Changes in the elasticity of learning-by-doing, ι_L, generate the second-largest response, but their quantitative impact remains small relative to that of α_G.

Results are especially insensitive to the strength of cross-sectoral spillovers. However, eliminating spillovers altogether leads to a further increase in the carbon-to-SCC ratio relative to the baseline. This observation is consistent with the idea that the goal to maintain some fossil knowledge dampens the planner's incentive to tax fossil energy as to avoid directing research to the green sector.

Variation in policy targets change the results only marginally. Higher fiscal pressure reduces the carbon tax wedge, in line with the view that directing learning through carbon pricing entails fiscal distortions. Imposing a net emissions target aligned with a 1.5°C pathway raises the carbon tax further above the SCC in the short run, while narrowing the wedge in the long run, relative to the baseline with a 2°C target.

Taken together, these results show that the optimal carbon tax is robustly above the SCC throughout the transition when learning-by-doing is endogenous, and that its magnitude is primarily governed by parameters that directly shape the strength of learning-by-doing in green production. Furthermore, mechanisms put forward in the analysis are confirmed.

Misallocation of innovation Panel B of Table 3 reports the implied misallocation of innovation as the average over the 2020-2074 horizon. Across robustness exercises, misallocation in the model with endogenous learning-by-doing remains negative, i.e. there is an underinvestment in fossil R&D. Misallocation measures appear tightly bounded, ranging between -3.2 and -4.2%. In all scenarios, misallocation gets amplified when learning-by-doing is endogenous relative to a model with exogenous expertise accumulation by between 3.0 and 8.5%.

The lowest levels of misallocation occur when government expenditure requirements are reduced, while the highest levels arise when fiscal pressure increases. This pattern highlights the role of fiscal constraints in shaping the trade-off between directing learning and innovative activity. The only outlier in the degree and direction of misallocation originates from the model without cross-sectoral spillovers. Here, the planner implements too little fossil-related innovation, on average by 33%. Hence, consistent with the argument put forward in this paper, cross-sectoral knowledge spillovers are pivotal on the social value of fossil-related innovation. The amplification of misallocation through learning-by-doing, conversely, remains close to the baseline model. Subsection F.1 investigates the effect of learning-by-doing on the misallocation of R&D in this model variant.

6 Conclusion

The transition to net-zero emissions demands a reorganization of production and innovation. While a transition of learning and innovation is central, much less is known about how these objectives should be balanced when climate policy is constrained by distortionary fiscal policy. To close this gap, the present paper studies optimal climate policy implementation in a dynamic general-equilibrium model with learning-by-doing, directed technical change, and distortionary taxation. The analysis highlights a tension between directing technical change and learning under fiscal distortions by deriving three quantitative results.

First, despite their fiscal inefficiency, optimal carbon taxes exceed the social cost of carbon in order to direct learning toward the green sector. Second, distortionary fiscal constraints induce a persistent misallocation of innovation as research subsidies necessitate generating government funds. Learning-by-doing aggravates the misallocation of innovation by motivating higher carbon taxation that directs research too quickly to green technologies. Third, a welfare decomposition shows that learning-by-doing amplifies the welfare costs of distortionary taxation, in particular when climate targets should be implemented. The main drivers are the fiscal distortions associated with promoting learning on green technologies.

References

Acemoglu, D. (2002). Directed technical change. The Review of Economic Studies, 69:781–809.

Acemoglu, D., Aghion, P., Barrage, L., and Hémous, D. (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. NBER Working Paper 31657, National Bureau of Economic Research.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hémous, D. (2012). The environment and directed technical change. American Economic Review, 102(1):131–166.

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. Journal of Political Economy, 124(1):52–104.

Aghion, P., Barrage, L., Donald, E., Hémous, D., and Liu, E. (2025). Green transitions along a supply chain. NBER Working Paper 33934, National Bureau of Economic Research.

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. Journal of Political Economy, 124(1):1–51.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. The Economic Journal, 99:116–131.

Barbieri, N., Marzucchi, A., and Rizzo, U. (2023). Green technologies, interdependencies, and policy. Journal of Environmental Economics and Management, 118:102791.

Barrage, L. (2020). Optimal dynamic carbon taxes in a climate-economy model with distortionary fiscal policy. Review of Economic Studies, 124(2):1–39.

Barwick, P. J., Kwon, H.-S., Li, S., and Zahur, N. (2025). Drive down the cost: Learning by doing and government policies in the global ev battery industry. NBER Working Paper 33378, National Bureau of Economic Research.

Belfiori, E. and Rezai, A. (2024). Implicit carbon prices: Making do with the taxes we have. Journal of Environmental Economics and Management, 125:102950.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? American Economic Review, 110(4):1104–1144.

Bollinger, B. and Gillingham, K. (2023). Learning-by-doing in solar photovoltaic installations. Mimeo, New York University and Yale University.

Bovenberg, A. L. and De Mooij, R. A. (1994). Environmental Levies and Distortionary Taxation. American Economic Review, 84(4):1085–1089.

Bovenberg, A. L. and Goulder, L. H. (1996). Optimal environmental taxation in the presence of other taxes: General-equilibrium analyses. American Economic Review, 86(4):985–1000.

Bovenberg, A. L. and Mooij, R. A. D. (1997). Environmental tax reform and endogenous growth. Journal of Public Economics, 63:207–237.

Bovenberg, A. L. and Smulders, S. (1995). Environmental quality and pollution-augmenting technological change in a two-sector endogenous growth model. Journal of Public Economics, 57(3):369–391.

Bovenberg, A. L. and Smulders, S. A. (1996). Transitional impacts of environmental policy in an endogenous growth model. International Economic Review, 37:861–893.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. American Economic Review, 101(3):471–475.

David, P. A. (1985). Clio and the economics of qwerty. The American Economic Review, 75(2):332–337.

Douenne, T., Hummel, A. J., and Pedroni, M. (2026). Optimal fiscal policy in a climate-economy model with heterogeneous households. The Economic Journal, page ueag006.

Edenhofer, O., Bauer, N., and Kriegler, E. (2005). The impact of technological change on climate protection and welfare: Insights from the model mind. Ecological Economics, 54:277–292.

European Commission (2020). Research and innovation for the european green deal. https://research-and-innovation.ec.europa.eu/strategy/strategy-research-and-innovation/environment-and-climate/european-green-deal_en. (Accessed on 09 February 2026).

Fischer, C. and Newell, R. G. (2008). Environmental and technology policies for climate mitigation. Journal of Environmental Economics and Management, 55(2):142–162.

Fried, S. (2018). Climate policy and innovation: A quantitative macroeconomic analysis. American Economic Journal: Macroeconomics, 10(1):90–118.

Fullerton, D. (1997). Environmental levies and distortionary taxation: Comment. The American Economic Review, 87(1):245–251.

Fullerton, D. and Kim, S. R. (2008). Environmental investment and policy with distortionary taxes, and endogenous growth. Journal of Environmental Economics and Management, 56:141–154.

Furman, J. L., Porter, M. E., and Stern, S. (2002). The determinants of national innovative capacity. Research Policy, 31:899–933.

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

Gerlagh, R., Kverndokk, S., and Rosendahl, K. E. (2014). The optimal time path of clean energy r&d policy when patents have finite lifetime. Journal of Environmental Economics and Management, 67:2–19.

Gerlagh, R., Zwaan, B. V. D., Hofkes, M. W., and Klaassen, G. (2004). Impacts of co 2-taxes in an economy with niche markets and learning-by-doing. Environmental and Resource Economics, 28:367–394.

Goulder, L. H. (1995). Environmental taxation and the double dividend: A reader's guide. International Tax and Public Finance, 2:157–183.

Goulder, L. H. and Mathai, K. (2000). Optimal co2 abatement in the presence of induced technological change. Journal of Environmental Economics and Management, 39:1–38.

Greaker, M., Heggedal, T. R., and Rosendahl, K. E. (2018). Environmental policy and the direction of technical change. Scandinavian Journal of Economics, 120:1100–1138.

Greiner, A. (2005). Fiscal policy in an endogenous growth model with public capital and pollution. The Japanese Economic Review, 56(1):67–84.

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(2):349–389.

Hassler, J., Krusell, P., and Smith, A. A. (2016). Environmental macroeconomics. In Handbook of Macroeconomics, volume 2, pages 1893–2008. Elsevier.

Hémous, D. (2016). The dynamic impact of unilateral environmental policies. Journal of International Economics, 103:80–95.

Hémous, D. and Olsen, M. (2021). Directed technical change in labor and environmental economics. Annual Review of Economics, 13:571–597.

Hettich, F. (1998). Growth effects of a revenue-neutral environmental tax reform. Zeitschrift for Nationalökonomie, 67:287–316.

Howell, S. T. (2017). Financing innovation: Evidence from r&d grants. American Economic Review, 107:1136–1164.

IPCC (2022). Summary for policymakers. In: *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [P.R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, J. Malley, (eds.)]. Cambridge University Press, Cambridge, UK and New York, NY, USA. doi: 10.1017/9781009157926.001.

IPCC (2023). Climate Change 2022 - Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.

Irwin, D. A. and Klenow, P. J. (1994). Learning-by-doing spillovers in the semiconductor industry. Journal of Political Economy, 102:1200–1227.

Jacobs, B. and van der Ploeg, F. (2019). Redistribution and pollution taxes with non-linear Engel curves. Journal of Environmental Economics and Management, 95:198–226.

Johnstone, N., Haščič, I., and Popp, D. (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. Environmental and Resource Economics, 45:133–155.

Jones, C. I. and Williams, J. C. (1998). Measuring the social return to r&d. The Quarterly Journal of Economics, 113(4):1119–1135.

Jones, L. E., Manuelli, R. E., and Rossi, P. E. (1993). Optimal taxation in models of endogenous growth. Journal of Political Economy, 101(3):485–517.

Jovanovic, B. and Nyarko, Y. (1996). Learning by doing and the choice of technology. Econometrica, 64:1299–1310.

Kalkuhl, M., Edenhofer, O., and Lessmann, K. (2012). Learning or lock-in: Optimal technology policies to support mitigation. Resource and Energy Economics, 34:1–23.

Kalkuhl, M., Edenhofer, O., and Lessmann, K. (2013). Renewable energy subsidies: Second-best policy or fatal aberration for mitigation? Resource and Energy Economics, 35:217–234.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles: I. The basic neoclassical model. Journal of Monetary Economics, 21:195–232.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. Quarterly Journal of Economics, 132(2):665–712.

Kruse-Andersen, P. K. (2023). Testing r&d-based endogenous growth models. Oxford Bulletin of Economics and Statistics, 85:1083–1110.

Kverndokk, S. and Rosendahl, K. E. (2007). Climate policies and learning by doing: Impacts and timing of technology subsidies. Resource and Energy Economics, 29:58–82.

Lemoine, D. (2024). Innovation-led transitions in energy supply. American Economic Journal: Macroeconomics, 16:29–65.

Makris, M. and Pavan, A. (2021). Taxation under learning by doing. Journal of Political Economy, 129(6):1878–1944.

Metcalf, G. E. (2003). Environmental levies and distortionary taxation: Pigou, taxation and pollution. Journal of Public Economics, 87:313–322.

Nemet, G. F. (2006). Beyond the learning curve: factors influencing cost reductions in photovoltaics. Energy Policy, 34:3218–3232.

Noailly, J. and Smeets, R. (2015). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. Journal of Environmental Economics and Management, 72:15–37.

Nykvist, B. and Nilsson, M. (2015). Rapidly falling costs of battery packs for electric vehicles. Nature Climate Change, 5:329–332.

OECD (2021). Average annual hours actually worked per worker. https://stats.oecd.org/Index.aspx?DataSetCode=ANHRS, (Accessed on 11 February 2021).

OECD (2022). Gross domestic product (GDP) (indicator). doi: 10.1787/dc2f7aec-en (Accessed on 03 September 2022).

Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. The Review of Economics and Statistics, 99(2):281–290.

Peterman, W. B. (2016). The effect of endogenous human capital accumulation on optimal taxation. Review of Economic Dynamics, 21:46–71.

Popp, D. (2002). Induced innovation and energy prices. American Economic Review, 92(1):160–180.

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

Popp, D. (2006). ENTICE-BR: The effects of backstop technology R&D on climate policy models. Energy Economics, 28:188–222.

Popp, D. (2016). Economic analysis of scientific publications and implications for energy research and development. Nature Energy, 1:16020.

Rezai, A. and van der Ploeg, F. (2017). Second-best renewable subsidies to de-carbonize the economy: Commitment and the green paradox. Environmental and Resource Economics, 66:409–434.

Robiou du Pont, Y., Jeffery, M. L., Gütschow, J., Rogelj, J., Christoff, P., and Meinshausen, M. (2017). Equitable mitigation to achieve the Paris Agreement goals. Nature Climate Change, 7(1):38–43.

Schneider, S. H. and Goulder, L. H. (1997). Achieving low-cost emissions targets. Nature, 389(6646):13–14.

United Nations Conference on Trade and Development (2022). $369 billion in investment incentives to address energy security and climate change. https://investmentpolicy.unctad.org/investment-policy-monitor/measures/4004/-369-billion-in-investment-incentives-to-address-energy-security-and-climate-change-. (Accessed on 09 February 2026).

United Nations Department of Economic and Social Affairs, Population Division (United Nations) (2022). World population prospects: The 2022 revision. https://population.un.org/dataportal/data/indicators/49/locations/900,840/start/2010/end/2100/table/pivotbylocation (Accessed on 04 August 2022).

United States Environmental Protection Agency (EPA) (2022). Trends in greenhouse gas emissions. https://nepis.epa.gov/Exe/ZyPDF.cgi/P1011XV3.PDF?Dockey=P1011XV3.PDF (Accessed on 02 February 2022).

US Energy Information Administration (EIA) (2023). Monthly Energy Review August 2023. https://www.eia.gov/totalenergy/data/monthly/, (Accessed on 06 September 2023).

Van den Bijgaart, I. (2017). The unilateral implementation of a sustainable growth path with directed technical change. European Economic Review, 91:305–327.

Van der Wijst, K., Byers, E., Riahi, K., Schaeffer, R., and Van Vuuren, D. (2023). Data for figure spm.5 of summary for policymakers of the synthesis report of the sixth assessment report, metadataworks.

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

Young, A. (1993). Invention and bounded learning by doing. Journal of Political Economy, 101(3):443–472.

A Model

This section spells out the equilibrium equations, simplifies household and government budgets, and discusses conditions for the existence and stability of a BGP.

A.1 Model equations

Household

Period utility $\log(C_t) - \chi \frac{H_t^{1+\sigma}}{1+\sigma}$

Budget $C_t = (1-\tau_{lt})w_t H_t + w_{st}S + \Pi_t + Gov_t^T - T_{xt}$

Optimality $C_t^{-1} = \mu_t p_t$

$\chi H_t^{\sigma} = \mu_t (1-\tau_{lt}) w_t$

where μ_t is the Lagrange multiplier on the budget constraint

Final good and Energy producers

Optimality $p_t \delta_y^{\frac{1}{\varepsilon_y}} Y_t^{\frac{1}{\varepsilon_y}} E_t^{-\frac{1}{\varepsilon_y}} = p_{Et}$

$p_t (1-\delta_y)^{\frac{1}{\varepsilon_y}} Y_t^{\frac{1}{\varepsilon_y}} N_t^{-\frac{1}{\varepsilon_y}} = p_{Nt}$

$p_{Et} E_t^{\frac{1}{\varepsilon_e}} F_t^{-\frac{1}{\varepsilon_e}} = p_{Ft} + \tau_{Ft}$

$p_{Et} E_t^{\frac{1}{\varepsilon_e}} G_t^{-\frac{1}{\varepsilon_e}} = p_{Gt}$

Definitions prices $p_t = \left[\delta_y p_{Et}^{1-\varepsilon_y} + (1-\delta_y) p_{Nt}^{1-\varepsilon_y}\right]^{\frac{1}{1-\varepsilon_y}}$

$p_{Et} = \left[(p_{Ft} + \tau_{Ft})^{1-\varepsilon_e} + p_{Gt}^{1-\varepsilon_e}\right]^{\frac{1}{1-\varepsilon_e}}$

Production $Y_t = \left(\delta_y^{\frac{1}{\varepsilon_y}} E_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1-\delta_y)^{\frac{1}{\varepsilon_y}} N_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}}\right)^{\frac{\varepsilon_y}{\varepsilon_y - 1}}$

$E_t = \left(F_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}} + G_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}}\right)^{\frac{\varepsilon_e}{\varepsilon_e - 1}}$

Intermediate good producers

Production $F_t = x_{Ft}^{\alpha_F} (A_{Ft} L_{Ft})^{1-\alpha_F}$

$N_t = x_{Nt}^{\alpha_N} (A_{Nt} L_{Nt})^{1-\alpha_N}$

$G_t = x_{Gt}^{\alpha_G} (A_{Gt} L_{Gt})^{1-\alpha_G}$

Labor demand $w_t = p_{Ft}^{\frac{1}{1-\alpha_F}} (1-\alpha_F) \alpha_F^{\frac{\alpha_F}{1-\alpha_F}} A_{Ft}$

$w_t = p_{Nt}^{\frac{1}{1-\alpha_N}} (1-\alpha_N) \alpha_N^{\frac{\alpha_N}{1-\alpha_N}} A_{Nt}$

$w_t = p_{Gt}^{\frac{1}{1-\alpha_G}} (1-\alpha_G) \alpha_G^{\frac{\alpha_G}{1-\alpha_G}} A_{Gt}$

Machine demand $x_{Fit} = (\alpha_F p_{Ft})^{\frac{1}{1-\alpha_F}} L_{Ft} A_{Fit}$

$x_{Nit} = (\alpha_N p_{Nt})^{\frac{1}{1-\alpha_N}} L_{Nt} A_{Nit}$

$x_{Git} = (\alpha_G p_{Gt})^{\frac{1}{1-\alpha_G}} L_{Gt} A_{Git}$

Learning $Q_{Ft} = (1-\delta_K) Q_{Ft-1} + F_t$

$$\begin{aligned}
& Q_{Gt} = (1-\delta_K)Q_{Gt-1} + G_t \\
\text{Productivity} \quad & A_{Fit} = \kappa Q_{Ft}^{\iota_L} K_{Fit}^{\iota_K} \\
& A_{Git} = \kappa Q_{Gt}^{\iota_L} K_{Git}^{\iota_K} \\
& A_{Nt} = (1+g_n)A_{Nt-1}
\end{aligned}$$

Machine producers

$$\begin{aligned}
\text{Price setting} \quad & p_{xFit} = \frac{1}{\alpha_F(1+\zeta_F)} \\
& p_{xNit} = \frac{1}{\alpha_N(1+\zeta_N)} \\
& p_{xGit} = \frac{1}{\alpha_G(1+\zeta_G)} \\
\text{Demand scientists} \quad & w_{st}(1-\tau_{sFt}) = \\
& \quad \eta\gamma(1-\alpha_F)\left(K_{Ft-1}^{1-\phi}K_{-Ft-1}^{\phi}\right)^{1-\nu}(s_{Fit})^{\eta-1}\frac{p_{Ft}F_t}{K_{Ft}} \\
& w_{st}(1-\tau_{sGt}) = \\
& \quad \eta\gamma(1-\alpha_G)\left(K_{Gt-1}^{1-\phi}K_{-Gt-1}^{\phi}\right)^{1-\nu}(s_{Git})^{\eta-1}\frac{p_{Gt}G_t}{K_{Gt}} \\
\text{Innovation} \quad & K_{Fit} = K_{Ft-1}(1-\delta_K) + \gamma(s_{Fit})^{\eta}\left(K_{Ft-1}^{1-\phi}K_{-Ft-1}^{\phi}\right)^{1-\nu} \\
& K_{Git} = K_{Gt-1}(1-\delta_K) + \gamma(s_{Git})^{\eta}\left(K_{Gt-1}^{1-\phi}K_{-Gt-1}^{\phi}\right)^{1-\nu}
\end{aligned}$$

Government

$$\begin{aligned}
& Gov_t^T = T_{Ft} + T_{Lt} - T_{Rt} \\
& \zeta_{Jt} = \frac{1-\alpha_J}{\alpha_J} \;\; for \;\; J \in \{F, G, N\}
\end{aligned}$$

Markets

$$\begin{aligned}
& H_t = L_{Ft} + L_{Nt} + L_{Gt} \\
& S = S_{Ft} + S_{Gt} \\
& C_t = Y_t - (X_{Ft} + X_{Gt} + X_{Nt})
\end{aligned}$$

A.2 Household and government budget

Government

$$\begin{aligned}
\text{Labor income taxation} \quad & T_{Lt} = \tau_{lt} w_t H_t \\
\text{Environmental policy} \quad & T_{Ft} = \tau_{Ft} F_t \\
\text{Research subsidization} \quad & T_{Rt} = w_{st}\tau_{sGt}S_{Gt} + w_{st}\tau_{sFt}S_{Ft} \\
\text{Monopoly correction} \quad & T_{xt} = \int_0^1 (p_{Fit}\zeta_t x_{Fit} + p_{Nit}\zeta_t x_{Nit} + p_{Git}\zeta_t x_{Git})\, di \\
\text{Government budget constraint} \quad & Gov_t^T = T_{Ft} + T_{Lt} - T_{Rt}
\end{aligned}$$

The household receives income from working, from engaging in science, from owning machine producing firms, and receiving transfers from the government:

Household

$$
\begin{aligned}
\text{Income} \quad & (1-\tau_{lt})w_t H_t + w_{st}S + \Pi_t + Gov_t^T - T_{xt} \\
\text{where} \quad & \Pi_t = \int_0^1 (\pi_{Fit} + \pi_{Git} + \pi_{Nit})\, di \\
\text{with} \quad & \pi_{Jit} = p_{Jit}(1+\zeta_{Jt})x_{Jit} - x_{Jit} - w_{st}(1-\tau_{Jt})s_{Jit} \quad \forall J, \quad s_{Nit} = 0 \;\; \forall t, \forall i
\end{aligned}
$$

Note that the household budget simplifies extensively as firm profits, income from science, and subsidies to firms and research cancel in equilibrium, that is:

$$
\begin{aligned}
& w_{st}S + \Pi_t = T_{xt} + T_{Rt} \\
\Leftrightarrow \quad & w_{st}(s_{Gt} + s_{Ft}) + \sum_J \int_0^1 (p_{Jit}(1+\zeta_{Jt})x_{Jit} - x_{Jit})di - w_{st}s_{Ft} - w_{st}(1-\tau_{sGt})s_{Gt} - w_{st}(1-\tau_{sFt})s_{Ft} \\
& = \int_0^1 (p_{Fit}\zeta_t x_{Fit} + p_{Nit}\zeta_t x_{Nit} + p_{Git}\zeta_t x_{Git})di + w_{st}\tau_{sGt}s_{Gt} + w_{st}\tau_{sFt}s_{Ft} \\
\Leftrightarrow \quad & 0 = 0.
\end{aligned}
$$

This observation is independent of the funding of research subsidies. Hence, the household budget becomes

$$
C = (1-\tau_{lt})w_t H_t + T_{Ft} + T_{Lt}. \tag{11}
$$

T_{Rt} remains relevant for the planning problem as it affects the government's budget constraint.

A.3 Stable interior balanced-growth path

The inherent difficulty in directed technical change frameworks lies in ensuring the existence and stability of interior balanced growth paths (BGP), where research occurs in both sectors. Two conditions must be satisfied. First, for an interior solution to exist, the marginal private gains from research in both sectors must be equalized. Second, for stability, the transitional dynamics must guide the economy toward the BGP rather than away from it.

As shown in Acemoglu (2002), the strength of cross-sectoral knowledge spillovers and the elasticity of substitution between input factors are central. Intuitively, when goods are close substitutes, even small technological advances in one sector can generate large shifts in demand, attracting more research and potentially triggering a reinforcing dynamic that leads to a corner solution, where all innovation is directed to a single sector. In contrast, when goods are more complementary, the demand response is muted, supporting the coexistence of innovation in both sectors. Importantly, what matters in general equilibrium for the stability is not only the elasticity of substitution between intermediaries, but the adjustment in input factors, i.e., labor input, that arises from changes in sectoral knowledge. In my model, under the assumption of equal capital shares, this elasticity is given by $\frac{\iota_k(1-\alpha)(1-\varepsilon)}{1-\eta}$.[45]

[45] I assume $\alpha_G = \alpha_F = \alpha$ for expositional reasons. The formula is derived in Appendix B.

Cross-sectoral knowledge spillovers offer another mechanism to sustain an interior BGP even if goods are substitutes: they allow the lagging sector to benefit from advances in the leading sector, dampening divergence. This is the approach chosen by Fried (2018). When such spillovers are weak, and when green and fossil goods are gross substitutes, the economy tends to converge to a corner solution, as in Acemoglu et al. (2012).

In my model, cross-sectoral spillovers are relatively strong, with a calibrated value of $\phi = 0.44$. High capital shares in the energy sector, $\alpha_F = 0.75$, and $\alpha_G = 0.88$, in combination with a relatively low elasticity of labor-augmenting productivity to knowledge, $\iota_K = 0.21$, and substantial decreasing returns to research, $\eta = 0.41$, further mitigate the reinforcing effect of knowledge advances. These features jointly support the existence of a stable interior BGP.

Learning-by-doing adds another layer of complexity. As a shift in the knowledge gap pushes production toward one sector, workers become more productive, which again raises the profits from research in that sector. Again, a low overall elasticity of output to growing expertise, $\frac{\iota_L(1-\alpha)(\varepsilon-1)}{1-\eta}$, diminishes the potential instability arising off the BGP.

B Derivations

This section derives the allocation of researchers as a function of sectoral expertise and knowledge stocks.

Intermediate sectors Demand for machines and labor follows from intermediate good producers' first order conditions. They maximize their profits given the production function $J = L^{1-\alpha} \int_0^1 A_{Jit}^{1-\alpha_J} x_{Jit}^{\alpha_J} di$ yielding demand for machines:

$$x_{Jt} = \left(\frac{p_{Jt}\alpha_J}{p_{xJt}} \right)^{\frac{1}{1-\alpha_J}} L_{Jt}A_{Jt}, \tag{12}$$

and labor demand

$$L_{Jt} = (1-\alpha_J)\frac{p_{Jt}}{w_t}J_t. \tag{13}$$

Eq. (12) shows that the elasticity of machine demand with respect to its price is $\frac{1}{1-\alpha_J}$. Substituting eq. (12) in the production function and setting the price of machines to its equilibrium price, $p_{xJt} = 1$, gives intermediate output at equilibrium machine input as

$$J_t = (p_{Jt}\alpha_J)^{\frac{\alpha_J}{1-\alpha_J}} L_{Jt}A_{Jt}. \tag{14}$$

Machine producers Machine producers maximize profits, eq. (1) in the text, by choosing the amount of researchers and the price of machines, taking into account demand for their machines, eq. (12). Machine producers set prices for machines to $p_{xJt} = 1$ under equilibrium subsidies

$$1 + \zeta_J = \frac{1}{\alpha_J}.$$

The subsidy ensures that in equilibrium the efficient amount of machines is produced. I drop the subscript i indicating machine producers for simplicity. The first order condition w.r.t. scientists reads

$$w_{st}(1-\tau_{sJt}) = \underbrace{(p_{xJt}(1+\zeta_{Jt})-1)}_{=\frac{1-\alpha_J}{\alpha_J}} \frac{\partial x_{Jt}}{\partial s_{Jt}}.$$

Using eq. (12), demand for machines responds to a marginal rise in s_{Jt} by

$$\frac{\partial x_{Jt}}{\partial s_{Jt}} = (p_{Jt}\alpha_J)^{\frac{1}{1-\alpha_J}} L_{Jt} \frac{\partial A_{Jt}}{\partial s_{Jt}}$$
$$\frac{\partial A_{Jt}}{\partial s_{Jt}} = \kappa \iota_K Q_{Jt}^{\iota_L} K_{Jt}^{\iota_K - 1} \frac{\partial K_{Jt}}{\partial s_{Jt}}$$

From the law of motion of innovation, eq. (2) in the text, the marginal innovation-product of scientists is:

$$\frac{\partial K_{Jt}}{\partial s_{Jt}} = \gamma\eta \left(K_{Jt-1}^{1-\phi} K_{-Jt-1}^{\phi}\right)^{1-\nu} s_{Jt}^{\eta-1}.$$

Assembling terms, the demand for scientists in equilibrium is given by

$$w_{st}(1-\tau_{sJt}) = \frac{(1-\alpha_J)}{\alpha_J} (\alpha_J p_{jt})^{\frac{1}{1-\alpha_J}} L_{Jt} \iota_K \frac{A_{Jt}}{K_{Jt}} \gamma\eta \left(K_{-Jt-1}^{\phi} K_{J_t-1}^{1-\phi}\right)^{1-\nu} s_{Jit}^{\eta-1}. \tag{15}$$

Equilibrium allocation of scientists The equilibrium allocation of scientists follows by dividing eq. (15) for sector G by sector F and rearranging terms:

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = \frac{\widehat{\alpha_G}}{\widehat{\alpha_F}} \frac{1-\tau_{sFt}}{1-\tau_{sGt}} \left(\frac{p_{Gt}^{\frac{1}{1-\alpha_G}}}{p_{Ft}^{\frac{1}{1-\alpha_F}}}\right) \frac{L_{Gt}}{L_{Ft}} \left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)} \left(\frac{A_{Gt}/K_{Gt}}{A_{Ft}/K_{Ft}}\right), \tag{16}$$

where $\widehat{\alpha}_J := \alpha_J^{\frac{\alpha_J}{1-\alpha_J}}(1-\alpha_J)$. Assuming equal capital shares, above equation reduces to eq. (8) in the text.

Relating sectoral expertise and the allocation of scientists To get at the general equilibrium effects of changing expertise on the direction of research, I first derive some helpful expressions. Combining labor demand, eq. (13), and production in equilibrium, eq. (14), uncovers an inverse relationship of productivity and prices:

$$p_{Jt}^{\frac{1}{1-\alpha_J}} = w_{Jt} \left(A_{Jt}\widehat{\alpha}_J\right)^{-1}.$$

Building the price ratio gives:

$$\frac{p_{Ft}^{\frac{1}{1-\alpha_F}}}{p_{Gt}^{\frac{1}{1-\alpha_G}}} = \frac{A_{Gt}\widehat{\alpha}_G}{A_{Ft}\widehat{\alpha}_F}. \tag{17}$$

Note that it follows from this relation of prices and productivity that the price effect and the marginal productivity effect cancel in eq. (16). Expertise affects the allocation of scientists through labor demand.

In general equilibrium, labor demand is determined by demand from energy producers for intermediates. Substituting sectoral output in labor demand, eq. (13), by energy producers' demand, $\frac{G_t}{F_t} = \left(\frac{p_{Ft}+\tau_{Ft}}{p_{Gt}}\right)^{\varepsilon_e}$, and taking the ratio yields the equilibrium allocation of labor as a function of prices:

$$\frac{L_{Gt}}{L_{Ft}} = (1+\widehat{\tau}_{Ft})^{\varepsilon_e} \left(\frac{p_{Ft}}{p_{Gt}}\right)^{\varepsilon_e-1} \frac{1-\alpha_G}{1-\alpha_F}. \tag{18}$$

where $\widehat{\tau}_{Ft} := \frac{\tau_{Ft}}{p_{Ft}}$. Eq. (18) captures the relationship of the relative market size, $\frac{L_{Gt}}{L_{Ft}}$, and prices. There are two counteracting effects of prices on labor demand. On the one hand, a higher price for fossil energy means a higher marginal product of fossil labor, raising demand for labor in the fossil sector. On the other hand, energy producers' demand shifts to green energy. The strength of the latter effect depends on the elasticity of substitution, ε_e. When the two goods are gross substitutes, $\varepsilon_e > 1$, the demand effect outweighs the price effect and a rise in the fossil price implies a reduction in fossil labor demand in equilibrium. The carbon tax only affects energy producers' demand leaving intermediate good producers' profits unaffected, ceteris paribus: independent of whether green and fossil goods are gross substitutes, the carbon tax unambiguously implies a direct shift towards green production.

Plugging in eq. (17) and eq. (18) in the equilibrium allocation of scientists, eq. (16), yields

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = \frac{1-\alpha_G}{1-\alpha_F}(1+\widehat{\tau}_{Ft})^{\varepsilon_e} \frac{1-\tau_{sFt}}{1-\tau_{sGt}} \left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)} \frac{K_{Ft}}{K_{Gt}} \left(w_t^{\alpha_G-\alpha_F} \frac{\widehat{\alpha_G}^{1-\alpha_G}}{\widehat{\alpha_F}^{1-\alpha_F}} \frac{A_{Gt}^{1-\alpha_G}}{A_{Ft}^{1-\alpha_F}}\right)^{\varepsilon_e-1},$$

where I use the assumption that machine producers are symmetric within sectors. Assuming equal capital shares and substituting the productivity function, eq. (4), we get eq. (9) in the text:

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = (1+\widehat{\tau}_{Ft})^{\varepsilon_e} \frac{1-\tau_{sFt}}{1-\tau_{sGt}} \left(\frac{Q_{Gt}^{\iota_L} K_{Gt}^{\iota_K}}{Q_{Ft}^{\iota_L} K_{Ft}^{\iota_K}}\right)^{(1-\alpha)(\varepsilon_e-1)} \frac{K_{Ft}}{K_{Gt}} \left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)}.$$

Strength of the expertise effect It is easy to see that the elasticity of the allocation of researchers to the gap in expertise is given by $\frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta}$. Four forces are relevant. First, the effect increases with the elasticity of productivity to expertise, ι_L, since advances in expertise translate more easily into higher sectoral productivity, price adjustments, and labor inputs. Second, the effect of the expertise gap attenuates by the rate $1-\alpha$ since labor demand only rises with labor-augmenting productivity with an elasticity of $1-\alpha$. The higher the labor share, $1-\alpha$, the stronger the effect of learning-by-doing on the allocation of researchers. Third, the tipping-on-toes effect reduces the elasticity. When it is strong, i.e., η is small, the elasticity declines because an increase in researchers within one sector raises the likelihood of duplicating ideas. As a result, market forces induce a more limited adjustment in research activity. Finally, the direction of the total effect of expertise on the allocation of researchers depends on the elasticity of substitution between green and fossil energy. On the one hand, higher sectoral expertise means more output, on the other hand, it implies a lower price. When the two goods are substitutes, $\varepsilon_e > 1$, the lower price for green goods fosters a rise in the ratio of green-to-fossil labor.

B.1 Learning-by-doing and the effect of carbon taxes on the allocation of research

Sketch of proof of Proposition 2 Consider the equilibrium allocation of researchers in eq. (9). For fixed knowledge stocks and research subsidies, the marginal effect of the carbon tax on the relative research allocation is strictly increasing in the expertise ratio Q_{Gt}/Q_{Ft}. Formally, defining

$$\Psi_t \ := \ \left(\frac{S_{Gt}}{S_{Ft}}\right),$$

differentiation with respect to $\widehat{\tau}_{F,t}$ yields

$$\frac{\partial \Psi_t}{\partial \widehat{\tau}_{F,t}} \ = \ \varepsilon_e(1+\widehat{\tau}_{F,t})^{\varepsilon_e-1}\left(\frac{Q_{Gt}}{Q_{Ft}}\right)^{\beta_Q} C_t, \qquad \beta_Q := \frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta} > 0,$$

where $C_t > 0$ collects all terms on the RHS in eq. (9) that do not depend on $\widehat{\tau}_{F,t}$ or Q_{Gt}/Q_{Ft}. Moreover,

$$\frac{\partial}{\partial(Q_{Gt}/Q_{Ft})}\left[\frac{\partial \Psi_t}{\partial \widehat{\tau}_{F,t}}\right] \ = \ \varepsilon_e(1+\widehat{\tau}_{F,t})^{\varepsilon_e-1}\beta_Q\left(\frac{Q_{Gt}}{Q_{Ft}}\right)^{\beta_Q-1} C_t \ > \ 0.$$

Hence, a higher, a-priori ratio of green-to-fossil expertise renders a marginal rise in the carbon tax more effective in directing research to the green sector, if $\beta_Q > 0$.

Sketch of proof of Proposition 3 To quantify how learning-by-doing alters the responsiveness of innovation to carbon taxation, rewrite eq. (9) in logs. Let $s_t := \log(S_{Gt}/S_{Ft})$, $k_t := \log(K_{Gt}/K_{Ft})$, and $q_t := \log(Q_{Gt}/Q_{Ft})$. Then eq. (9) implies

$$s_t = \frac{\varepsilon_e}{1-\eta}\log(1+\widehat{\tau}_{F,t}) + \frac{1}{1-\eta}\log\left(\frac{1-\tau_{sF,t}}{1-\tau_{sG,t}}\right) + \beta_K\, k_t + \beta_Q\, q_t + \beta_{\text{spill}}\, k_{t-1}, \tag{19}$$

where

$$\beta_K := \frac{\iota_K(1-\alpha)(\varepsilon_e-1)-1}{1-\eta}, \qquad \beta_{\text{spill}} := \frac{(1-2\phi)(1-\nu)}{1-\eta}.$$

Only k_t and q_t depend on the current allocation s_t, while k_{t-1} is predetermined. Define the equilibrium condition $\Phi(s_t, \widehat{\tau}_{F,t}) = 0$ by subtracting the RHS of eq. (19) from s_t. An application of the Implicit Function Theorem yields

$$\frac{ds_t}{d\widehat{\tau}_{F,t}} = -\frac{\Phi_{\widehat{\tau}}}{\Phi_s} = \frac{\dfrac{\varepsilon_e}{(1-\eta)(1+\widehat{\tau}_{F,t})} + \beta_Q\dfrac{\partial q_t}{\partial \widehat{\tau}_{F,t}}}{1-\beta_K\dfrac{\partial k_t}{\partial s_t} - \beta_Q\dfrac{\partial q_t}{\partial s_t}}. \tag{20}$$

C Social planner

The solution to the social planner's problem is defined as an allocation $\{L_{Ft}, L_{Gt}, L_{Nt}, X_{Ft}, X_{Gt}, X_{Nt}, C_t, H_t, S_{Ft}, S_{Gt}, S_{Nt}\}$ for each period which maximizes the social welfare function

$$\max_{\{C_t,H_t,L_{Jt},S_{Jt},X_{Jt}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_t, H_t)$$

$$\begin{aligned} \text{s.t.} \quad & \omega F_t - \delta \leq \Omega_t, \\ & C_t + X_{Ft} + X_{Gt} + X_{Nt} = Y_t, \\ & L_{Ft} + L_{Gt} + L_{Nt} \leq H_t, \\ & S_{Ft} + S_{Gt} \leq S, \\ & K_{Jt+1} = (1-\delta_K)K_{Jt} + \Phi(S_{Jt}, K_{Jt}, K_{-Jt}), \qquad J \in \{F, G\}, \\ & Q_{Jt+1} = (1-\delta_K)Q_{Jt} + J_t, \qquad J \in \{F, G\}, \\ & A_{Nt+1} = (1+\gamma_n)A_{Nt}, \\ & K_{J0}, Q_{J0}, A_{J0} \text{ given.} \end{aligned}$$

Production of Y_t is defined by the equations describing production. It holds that $X_{Jt} = \int_0^1 (x_{Jit})\, di$ and $S_{Jt} = \int_0^1 (s_{Jit})\, di$. Φ stands in for the production of new innovation.

D Numerical appendix

Since I cannot solve explicitly for the optimal policy over an infinite horizon, I approximate the problem by truncating it after period T and approximating future welfare using a constant growth-rate approximation. The infinite-horizon problem is approximated by truncating the planner's problem at $T = 20$, corresponding to a 100-year horizon, and appending a continuation value. More formally, the planner's objective function becomes:

$$\sum_{t=0}^{T} \beta^t u(C_t, H_t) + PV.$$

I define the continuation value, PV, as the period utility over the infinite horizon starting from the last explicit maximization period:

$$PV = \sum_{s=T+1}^{\infty} \beta^s u(C_s, H_s).$$

I make two simplifying assumptions to derive the continuation value. First, I assume that the consumption share, c_s, with $C_s = c_s Y_s$, is constant from period $T+1$ onward. As a result, consumption grows at the same rate as output. For the purpose of evaluating the continuation value, I approximate this growth rate by the model's calibrated balanced-growth-path growth rate, which is pinned down by exogenous productivity growth in the non-energy sector.[46] Under these assumptions, I can rewrite future consumption as $C_s = (1+\gamma_n)^{s-T} C_T \, \forall s \geq T+1$. Second, hours worked remain at their value in the last explicit optimization period: $H_s = H_T \, \forall s \geq T+1$.

[46]Note that I do not assume that the post-T economy is optimally described by the continuation value. Rather, the continuation value serves as a parsimonious approximation to the infinite-horizon utility associated with the technology and expertise levels inherited at period T.

Given the per-period utility function

$$u(C_s, H_s) = \log C_s - \chi \frac{H_s^{1+\sigma}}{1+\sigma},$$

and using $C_s = (1+\gamma_n)^{s-T} C_T$ and $H_s = H_T$ for all $s \geq T+1$, period-s utility can be written as

$$u(C_s, H_s) = \log C_T + (s-T)\log(1+\gamma_n) - \chi \frac{H_T^{1+\sigma}}{1+\sigma}.$$

Substituting this expression into the definition of the continuation value yields

$$PV = \sum_{s=T+1}^{\infty} \beta^s \left(\log C_T + (s-T)\log(1+\gamma_n) - \chi \frac{H_T^{1+\sigma}}{1+\sigma} \right).$$

Evaluating the (arithmetico-)geometric sums in closed form gives

$$PV = \beta^{T+1} \left(\frac{\log C_T}{1-\beta} + \frac{\log(1+\gamma_n)}{(1-\beta)^2} - \frac{\chi}{1-\beta} \frac{H_T^{1+\sigma}}{1+\sigma} \right).$$

This approach differs from terminal closures commonly used in the literature, where balanced-growth or steady-state conditions are imposed to close infinite-horizon planning problems(e.g. Barrage, 2020; Jones et al., 1993). Such a closure is not appropriate in the present setting. The green transition is characterized by prolonged, policy-driven adjustments in key sectoral ratio which do not satisfy BGP conditions. I therefore approximate the infinite-horizon problem by appending a continuation value rather than by imposing steady-state policy or allocation conditions.

E Calibration

This section details the knowledge stock measurement and the calibration of the productivity function.

E.1 Measuring knowledge stocks from patent data

The distribution of initial knowledge stocks is a crucial driver of the optimal environmental policy, as it determines the relative productivity of researchers across sectors through knowledge spillovers. To measure sectoral knowledge stocks, I use the universe of patents granted by the United States Patent and Trademark Office (USPTO) and filed by U.S. applicants[47] from the European Patent Office (EPO)'s database PATSTAT. I consider patents filed between 1950 to 2014. The number of granted patents displays a sharp reduction after this year due to the time which elapses from applying for protection to a patent being granted. The data on granted patents for more recent years is, therefore, incomplete.

[47]These may be companies, individuals, or the government. I include government patent applicants because the innovation remains important for knowledge spillovers. Such patents account only for 1.5% of all patents considered.

To classify patents into the distinct sectors of the model, I rely on the classification provided by a joint work of the International Energy Agency and the EPO.[48].

Using the number of patents related to green, fossil, and non-energy[49] technologies, I calculate a measure of the knowledge stock within sectors based on the *perpetual inventory method* which assumes that knowledge accumulates over time and depreciates:

$$K_{Jt} = (1 - \delta_K)K_{Jt-1} + R_{Jt},$$

where R_{Jt} stands in for new patents in sector J at time t. Depreciation of knowledge captures that knowledge becomes obsolete overtime as it is overrun by new innovation. To achieve consistency with the model, one period in the perpetual inventory model is set to 5 years.

Figure 5 depicts the evolution of the annual knowledge stock by sectors over time. The fossil-related knowledge stock exceeds green knowledge, albeit a catching up of green knowledge in the mid-2000's, the stock of fossil knowledge remained higher. In recent years, patenting in the energy sector reduced, and depreciation of knowledge caused a reduction of the knowledge stock in the green and the fossil sector. However, this drop is stronger in the green sector. As a result, the gap between fossil and green knowledge stocks widened in the late 2010's. The figure also highlights the drop in granted patents since the mid-2010s owing to the duration of the application process.

E.2 Measuring sectoral expertise

First, I derive scale-dependent measures of expertise based on historic cumulative output as described by the model's LOM for expertise, eq. 3. I take sector-specific output from (EIA, 2023, Table 1.1), defining nuclear and renewables as the model's *green* energy and fossil fuels as *fossil* energy. These series provide a proxy for the ratio of green-to-fossil expertise. Jointly with the long-run ratio of knowledge stocks they provide the ratio of fossil-to-green productivities, A_F/A_G, on the BGP. Based on this information and the additional calibration steps that determine the non-research side of the model, I derive base-year values for fossil and green output in the model. Importantly, these values are independent of base-year levels of expertise. Second, I employ these model-implied output levels to construct time series of fossil and green energy production using historic sector-specific growth rates from the data going back until 1950. I then re-calculate model-consistent expertise as cumulative output based on these synthetic time series.

E.3 Calibration of learning rates

To calibrate the elasticity of labor-augmenting productivity, A_J, in the model with respect to expertise, I proceed as follows. The literature commonly estimates learning rates as the effect of the doubling of production on total factor productivity (TFP). The elasticity in my model is linked to these estimates in two steps. First, TFP in sector J in the model is given by, $A_J^{1-\alpha_J}$. Define $\iota_{\tilde{L}J}$ and $\iota_{\tilde{K}J}$ as $A_J^{1-\alpha_J} = Q_{Jt}^{\iota_{\tilde{L}J}} K_{Jt}^{\iota_{\tilde{K}J}}$.

[48]The table of classifications of green technologies can be found here: https://link.epo.org/web/patents_and_the_energy_transition_study_en.pdf. The equivalent table for fossil-based technologies is given here: https://link.epo.org/web/patents_and_the_energy_transition_study_annex_en.pdf

[49]Non-energy patents are determined residually but not used in the model

Figure 5: Annual knowledge stock by sector in number of patents

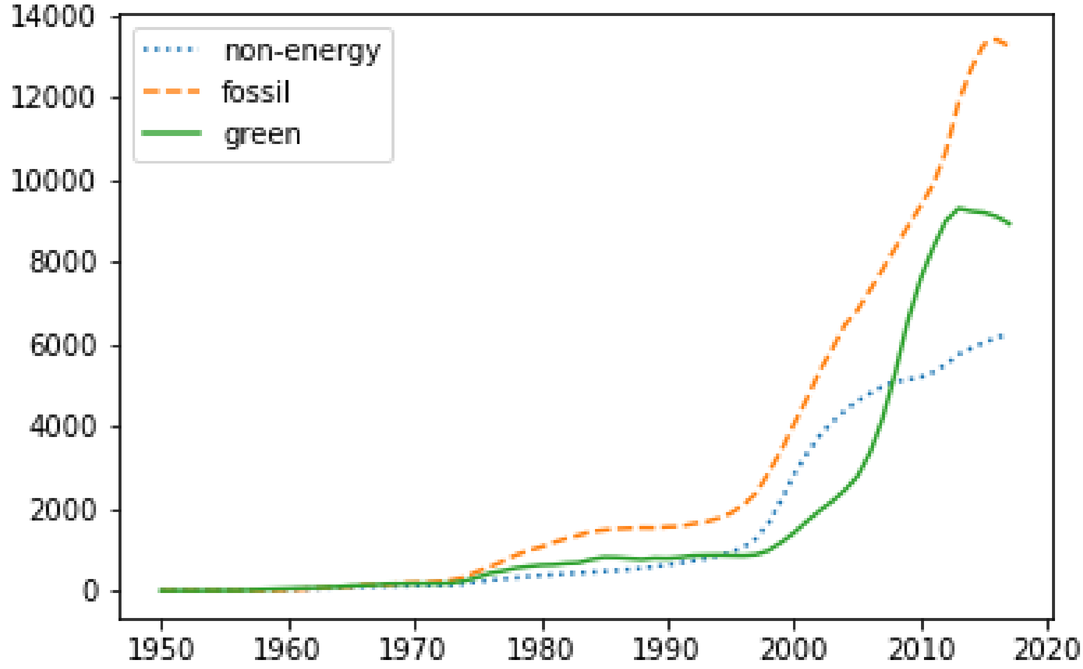

Notes: Knowledge stock per research process by sector in the US based on the number of granted patents by the US patent authority (USPTO). Only patents filed by U.S. applicants are considered, the respective filing date is shown on the x-axis. Data comes from the EPO's patent data bank PATSTAT. I classify patents by sector based on definitions derived by patent and energy experts as described here: https://link.epo.org/web/patents_and_the_energy_transition_study_en.pdf.

Sector-specific capital shares at sector-unspecific $\tilde{\iota_L}$ would imply sector-specific elasticities of labor productivity to expertise. Since I do not want to allow for sector-specific learning rates, I allow the elasticities of total factor productivity to vary across sectors. This means that within a given sector, workers learn equally quickly how to use new technologies. However, sectors with a lower labor share experience a smaller effect on TFP as their experience or the technology stock increases.

These parameters match the elasticity of TFP and learning in the data. Second, the literature provides estimates for learning rates, LR, which are commonly defined as the effect of a doubling in cumulative output as a proxy for workers' expertise or the technology stock on prices or TFP, hence $LR = \frac{TFP(2Q)-TFP(Q)}{TFP(Q)}$ with the function $TFP()$ relating TFP to expertise, Q. Using the functional forms in my model and solving for the elasticity yields: $\iota_{\tilde{L}J} = \frac{\log(LR+1)}{\log(2)}$.

I match the elasticity of labor-augmenting productivity with respect to expertise based on the evidence in (Irwin and Klenow, 1994). In a study relating prices and cumulative output in the semiconductor industry, they find learning rates of 20%. I match their estimate to the non-energy sector[50], since semiconductors are used in a variety of sectors

[50]Even though growth in the non-energy sector in the model is exogneous, I can match the hypothetic learning rate in the non-energy sector to the data which only requires knowledge of the labor share in the non-energy sector.

not necessarily related to energy. $\iota_{\tilde{LN}}$ is then backed out as $\iota_{\tilde{LN}} = -\frac{ln(1-0.2)}{ln(2)} = 0.32$.[51] Assuming that elasticities of labor-augmenting productivities are the same across sectors, we have that $\iota_L = \frac{\iota_{\tilde{LN}}}{1-\alpha_n} = 0.503$.[52]

Furman et al. (2002) quantify the effect of increases in the patent stock on aggregate productivity. A doubling of the patent stock relates roughly to a 10% increase in TFP. I match this measure to the non-energy sector yielding $\iota_K = 0.2148$.[53] All in all, the curvature of labor-augmenting productivity to technology and expertise is below unity establishing a concave relationship and decreasing returns to the composite of knowledge and expertise.

F Additional Results

F.1 Cross-sectoral spillovers

Even in the absence of cross-sectoral knowledge spillovers, learning-by-doing amplifies the misallocation of researchers (see Figure 6a). The underlying mechanism, however, differs from the baseline case. When spillovers are shut down, fossil research no longer generates indirect social value through future green innovation. Instead, the planner fosters a more rapid shift in R&D to green innovation, and carbon taxation pushes both, expertise and innovation, in a desired direction.

Figure 6b depicts the isolated effects of the change in the carbon tax, in expertise and in research subsidies when learning is endogenous relative to the exogenous benchmark. The rise in the carbon tax once learning on green technologies requires policy support, helps lower the ratio of fossil-to-green R&D in line with the government's target by between 2.2 and 2.8%.

Looking at the policy response when lump-sum finance for research subsidies is available, is informative on how research subsidies should be adjusted. Now, when learning-by-doing is endogenous, the planner seeks to spur green research even more and in addition to the rising carbon tax; compare the thin blue-dashed line.

This observation can be explained by Proposition 2, which says that the marginal effect of the carbon tax on the ratio of green-to-fossil R&D positively depends on the green-to-fossil expertise gap. As green expertise reduces in the endogenous scenario, the carbon tax is less effective in directing research. The subsidy on green research compensates this.

When research subsidies have to be financed with distortionary instruments, however, the use of subsidies is limited; compare the thick blue-solid graph. As a result, the misallocation of researchers widens relative to the scenario with exogenous expertise.

[51] Note the inverse relationship of unit costs, C, and TFP: $C = \frac{1}{A^{1-\alpha_n}} = Q^{-\iota_{\tilde{L}}} K^{-\iota_{\tilde{L}}}$.

[52] This results in the following elasticities of fossil and green TFP with respect to expertise: $\iota_{\tilde{LF}} = 0.1258$ and $\iota_{\tilde{LG}} = 0.0706$. Owing to the lower labor share in the green sector, a percentage increase in labor-augmenting productivity in the green sector has a smaller effect on green TFP than in the fossil sector.

[53] Sector-specific elasticities of TFP are given by $\iota_{\tilde{KG}} = 0.0302$,$\iota_{\tilde{KF}} = 0.0537$.

Figure 6: Misallocation of innovation under fiscal distortions: Absent cross-sectoral spillovers

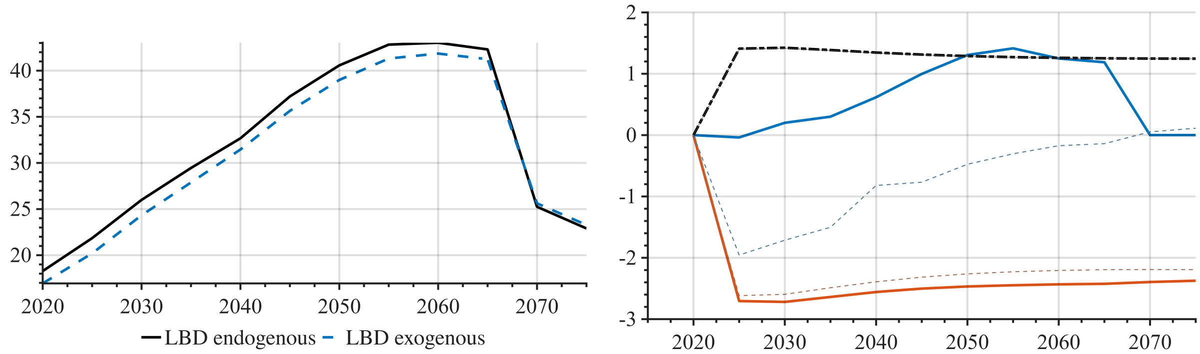

(a) Misallocation: Deviation from allocation under lump-sum–financed subsidies in %

(b) Decomposition: Effect of endogenizing learning-by-doing on fossil-to-green R&D ratio

Notes: The figure depicts results in a model where cross-sectoral spillovers are switched off, i.e., $\phi = 0$. Panel 6a quantifies the misallocation of innovation induced by fiscal distortions, measured as the percentage deviation of the fossil-to-green R&D ratio in the fully distortionary regime from a benchmark with lump-sum financed research subsidies for the model with and without endogenous learning-by-doing. Finally, Panel 6b decomposes the difference in the logged ratio of fossil-to-green research allocation between the model with endogenous and exogenous learning-by-doing into ceteris paribus contributions from carbon taxes, research subsidies, and expertise. Reported percentages are computed as $(\exp(x) - 1) \times 100$ and measure changes in the fossil-to-green research ratio relative to the exogenous-learning benchmark. The x-axis reports the first year of the five-year period to which each variable refers.

Table 3: Robustness: Carbon tax and misallocation of innovation

Panel A: Ratio of carbon tax to SCC (by period)

Period	2020-2024		2045-2049		2070-2074	
Parameter	LBD	Exo. LBD	LBD	Exo. LBD	LBD	Exo. LBD
Baseline	1.0354	0.9887	1.0259	0.9828	1.0227	0.9784
Lower α_G	1.3514	1.0098	1.1288	0.9879	1.1065	0.9834
Higher α_G	1.0095	0.9876	1.0083	0.9822	1.0058	0.9769
Lower ι_L	1.0125	0.9882	1.0049	0.9828	1.0007	0.9784
Higher ι_L	1.0558	0.9891	1.0450	0.9828	1.0438	0.9783
No cross-sectoral spillovers	1.0420	0.9960	1.0304	0.9888	1.0252	0.9824
Lower ϕ	1.0355	0.9887	1.0260	0.9829	1.0228	0.9785
Higher ϕ	1.0354	0.9886	1.0259	0.9828	1.0227	0.9783
Lower η	1.0368	0.9897	1.0268	0.9834	1.0236	0.9789
Higher η	1.0335	0.9872	1.0244	0.9820	1.0213	0.9776
Lower Gov_t^{min}	1.0392	0.9910	1.0304	0.9868	1.0282	0.9832
Higher Gov_t^{min}	1.0313	0.9860	1.0207	0.9783	1.0164	0.9727
Stricter emissions target	1.0381	0.9887	1.0277	0.9812	1.0219	0.9784

Panel B: Misallocation of innovation and amplification by LBD (time averages)

Parameter	Misallocation of innovation ($\frac{S_F}{S_G}$)	Effect of LBD in %
Baseline	−3.6798	4.3452
Lower α_G	−4.0880	4.0341
Higher α_G	−3.5228	3.8991
Lower ι_L	−3.5548	2.9494
Higher ι_L	−3.8712	7.1106
No cross-sectoral spillovers	32.6750	4.4011
Lower ϕ	−3.3797	5.1123
Higher ϕ	−3.8849	4.5436
Lower η	−3.7571	8.5456
Higher η	−3.5945	4.3633
Lower Gov_t^{min}	−3.1784	6.3342
Higher Gov_t^{min}	−4.2036	3.9501
Stricter emissions target	−3.8635	4.9111

Notes: Panel A reports the ratio of the optimal carbon tax to the SCC. Panel B reports the percentage misallocation of researchers relative to the allocation under lump-sum financed subsidies. All rows vary the indicated parameter individually relative to the baseline calibration. *Lower* α_G and *Higher* α_G indicate that the green labor share is fixed at $\alpha_G = 0.91$ following Fried (2018), and $\alpha_G = 0.75$. *No cross-sectoral spill.* refers to the model with $\phi = 0$, *Lower* ϕ and *Higher* ϕ refer to $\phi = 0.31$ similar to Aghion et al. (2016) and $\phi = 0.60$ as in Fried (2018). *Lower* ι_L (*Higher* ι_L) indicate minus (plus) 50% relative to the baseline level $\iota_L = 0.50$. *Lower* η and *Higher* η refer to 0.19 and 0.79 the values used in Hart (2019) and Fried (2018), respectively. *Lower* Gov_t^{min} (*Higher* Gov_t^{min}) refers to a 10% decrease (increase) in government expenditures in each period relative to the baseline path. Finally, the *Stricter emissions target* is consistent with a 1.5°C warming based on IPCC (2022).