Optimal Climate Policy under Distortionary Fiscal Constraints: The Innovation-Learning Dilemma

Sonja Dobkowitz *

March 26, 2026

Abstract I study optimal climate policy in a model with distortionary fiscal constraints, learning-by-doing, and directed technical change. The framework gives rise to an innovation–learning dilemma, in which policies that promote learning and those that direct innovation compete for limited fiscal resources. In a calibration to the U.S., I show that the optimal carbon tax exceeds the social cost of carbon to steer learning toward green technologies. Fiscal distortions shift the socially desirable allocation of research and induce a persistent misallocation that is amplified by learning-by-doing. A welfare decomposition shows that learning-by-doing increases the welfare losses associated with both fiscal distortions and the misallocation of research, corroborating the presence of an innovation–learning dilemma under distortionary fiscal constraints.

JEL classification: H21, H23, O3, Q54, Q55

* 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, Christian Bayer, 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, Katrin Millock, Hélène Ollivier, 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 shift in how economies produce. Governments worldwide rely on research subsidies to support this shift,[1] and a growing literature shows that such subsidies help accelerate the transition.[2] Yet financing subsidies necessitates the use of distortionary fiscal instruments to raise revenues—a constraint largely ignored in the literature on directed technical change and the green transition.

At the same time, economies must build expertise to adopt green technologies, and learning-by-doing plays a central role in driving this process.[3] But distortionary fiscal policies reduce labor supply, undermining the accumulation of experience that is critical for the green transition. Taken together, governments face a trade-off: directing innovation through research subsidies requires fiscal resources, yet the very instruments used to raise them can suppress the labor input essential for effective learning.

To analyze the implications of this trade-off for the optimal policy, 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.[4] It anticipates that net emissions must be limited in the near term and

[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 theoretic 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 research by financially constraint firms, attesting an important role for green energy technologies. Other contributions are Johnstone et al. (2010) and Popp (2016).

[3] Barwick et al. (2025) find a substantial role of learning-by-doing for the cost reduction of electric vehicles. Irwin and Klenow (1994) and Nykvist and Nilsson (2015) are other examples for the importance of learning-by-doing. Smaller effects are found in Bollinger and Gillingham (2023) and Nemet (2006).

[4] I abstract from targeted subsidies for learning-by-doing, 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

eliminated entirely in the long run.[5] Calibrating the model to the U.S. economy, I characterize the optimal policy mix during the transition toward net-zero emissions and thereafter.

Before turning to the quantitative analysis, I use the model to establish a general-equilibrium relationship between sectoral expertise and the allocation of research. I show that R&D is directed toward the sector with more experienced workers, because new technologies can be used more productively when expertise is high. Through this channel, fiscal distortions dynamically alter the *desired* research allocation the planner would like to implement during the transition to net-zero emissions, relative to the first best. In particular, under fiscal distortions the emission target is implemented at a higher fossil-to-green energy ratio, maintaining fossil expertise, and fossil-related research remains socially valuable.

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 taxes.[6] By contrast, public economics analyses emphasize that learning-by-doing can further raise the distortionary costs of taxation 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 about 3% above the social

between learning and optimal fiscal tax design. Moreover, targeted learning subsidies would not eliminate the fiscal trade-offs at the core of the analysis, as they too would require distortionary financing.

[5] In so doing, my paper contributes a cost-effectiveness analysis in a directed technical change framework. 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).

[6] 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. (2022). 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.

cost of carbon throughout the transition.[7] By shifting production toward low-carbon technologies, 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 much 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 fiscal cost of carbon taxation and strengthens the case for using it to direct learning.

Second, I study how fiscal distortions affect the allocation of research. I distinguish between the *desired* allocation of research—the allocation the planner would implement with lump-sum–financed research subsidies—and the allocation that can be implemented under distortionary finance, when research subsidies are fiscally costly. I define the gap between the two as the misallocation of researchers due to fiscal distortions.

Fiscal distortions alter the desired trajectory of innovation. Under distortionary finance, the emissions target is implemented at a smaller scale, allowing for a higher ratio of fossil energy in production. This slows the accumulation of green expertise relative to the first best. Through the expertise effect established analytically, the desired allocation of research, therefore, features a persistently higher share of fossil-related R&D. What matters for this result is not the overall level of fossil energy use—which is fixed—but the composition of energy production.

I document a persistent misallocation of research effort under distortionary finance. As research subsidies must be financed with distortionary taxes, the planner limits their use and tolerates a composition of research that deviates from the benchmark with lump-sum financed subsidies. I find that the implemented allocation features too low a share of fossil-related R&D. The benefit of maintained investment in fossil-based R&D crucially hinges on the existence of cross-sectoral knowledge spillovers. They raise the social value of fossil-related research which can stimulate future ideas on green technology.[8]

[7] The social cost of carbon is defined as the shadow price of the emissions constraint; in a first-best environment with lump-sum taxation and instruments to directly target sector-specific learning, the carbon tax would coincide with this value.

[8] Barbieri et al. (2023) provide empirical evidence for the importance of knowledge spillovers from non-green to green knowledge.

Endogenous learning-by-doing amplifies the misallocation of research. To foster green learning-by-doing, the planner optimally raises the carbon tax, which directly tilts private research incentives toward green technologies. In principle, a sufficiently strong increase in subsidies to fossil research could offset this side effect of carbon taxation. In practice, however, the need to finance research subsidies with distortionary instruments weakens their response. As a result, the implemented allocation features too little fossil-related R&D relative to the desired allocation, and the misallocation of researchers widens.

Third, I quantify the welfare costs of fiscal distortions using consumption-equivalent variation (CEV), defined as the percentage of per-period consumption the representative household would be willing to forgo to avoid distortionary taxation. I show that learning-by-doing amplifies the welfare cost of fiscal distortions by 15%. This occurs for two reasons. First, even absent an emission target, fiscal distortions are more costly when learning-by-doing is endogenous because they reduce dynamic productivity gains through lower expertise accumulation. Second, when an emission target is imposed, welfare losses increase further because learning-by-doing strengthens the case for relying on carbon taxes—which are relatively distortionary. Finally, the welfare gains from correcting the misallocation of research are larger when learning-by-doing is endogenous, consistent with the presence of an innovation–learning dilemma under fiscal constraints.

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 into this framework, 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 revenues. 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. Conversely, research subsidies become more important in the long run.

Moreover, I show that fiscal constraints change the beneficial allocation of researchers. As the emission target is implemented at a lower scale and higher share of fossil energy, the desired allocation of research efforts consists of a higher maintained share of fossil-related R&D relative to first best. This mechanism is absent from existing directed innovation models, where fiscal finance does not constrain policy design.[9]

A second, more general contribution to the directed technical change literature is to show how learning-by-doing dynamically 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. I extend this logic to the direction of innovation: research is endogenously directed toward the sector with greater accumulated experience. The mechanism introduces a channel through which fiscal distortions matter for the optimal allocation of researchers during the transition to net-zero emissions.

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

[9]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); Zwaan et al. (2002); Popp (2004, 2006); Fischer and Newell (2008); Gerlagh et al. (2009, 2014). All these papers abstract from fiscal constraints and study models different from the modern directed technical change literature.

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 fiscal constraints in environmental policy. 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 et al., 2022). 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 set-up. In particular, my findings point to the additional gains of carbon taxes that outweigh the fiscal costs which commonly lower their level below the social cost of carbon.

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 steered to meet a predetermined emissions target.

Outline The remainder of the paper is structured as follows. Section 2 presents the model. I derive theoretic results in Section 3 and elaborate on the calibration in Section 4. The quantitative results are discussed in Section 5. 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. Together, these ingredients create an innovation-learning dilemma.

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[10]. It also supplies scientific work in a fixed amount: S.[11] The household behaves according to solving below problem each period:

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

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 $\tau_{\iota t}$.[12] I

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

[11] It is common to fix the supply of scientists in the literature on directed technical change in order to simplify the analysis (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).

[12] Since this paper focuses on the effects of labor income taxes through their effect on the level of production, redistributive and heterogeneous effects of non-linear labor taxes can be neglected. A recent literature studies the effects of labor taxation and redistribution on the direction of research through a labor supply, or market size, effect with heterogeneous agents. See for instance, Loebbing (2019).

use a utility specification following King et al. (1988) that admits balanced growth with constant hours worked. The parameter σ captures the inverse Frisch elasticity. T_{xt} indicates lump-sum transfers to finance subsidies for the correction of monopolistic competition.[13]

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:

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

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 will 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}.

[13] Allowing for such instruments to exists 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. These transfers are given by $T_{xt} = \sum_J \int_0^1 (p_{xJit} \zeta_{Jt} x_{Jit}) di$.

Capital shares, α_J, are sector specific. Intermediate good producers maximize profits:

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

where w_{lt} is the wage rate for labor, which is equal across all sectors due to the assumption of free labor mobility, and p_{xJit} denotes the price of machines from producer i in sector J.

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, machine producers in the fossil and green sector decide on the amount of scientists to employ. Productivity in sector N evolves according to $A_{Nt} = (1 + g_n)A_{Nt-1}$ and the amount of scientists is fixed $s_{Nit} = 0$ for all firms and periods.

Demand for machines increases with their productivity which again is a function of technological progress. This provides the incentive to invest in research in the green and fossil sector. 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.[14] 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.[15]

In what follows I focus on the green and fossil sector. 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

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

[15]I introduce this policy to abstract from market imperfections as a driver of the results. The so-financed subsidy remains in place also in the distortionary setting.

normalization of one subsidy innocuous. However, when fiscal distortions are present, such normalization has real implications: it implicitly determines which sector's subsidy is fiscally more costly. 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 researcher productivity depends on the existing stock of knowledge. I use the terms *knowledge* and *technology* to distinguish between innovation as an input into idea generation and as an input into production, respectively. *Expertise* refers to workers' ability to operate technologies efficiently and evolves through learning-by-doing.

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

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

Here, γ captures the productivity of researchers, and δ_K denotes the depreciation rate of knowledge. The parameter $\eta < 1$ introduces diminishing returns to research effort, capturing a "stepping-on-toes" effect: as more researchers work on the same sector, the likelihood of duplicating ideas increases. A more equal allocation of researchers across sectors alleviates the duplication 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, captured by $K_{-J,t-1}^{\phi(1-\nu)}$. These spillovers imply that knowledge generated in one sector enhances the productivity of research in the other. 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 governed by $\nu \in (0,1)$. 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).[16]

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}. Consistent with the free mobility of labor, workers' expertise is sector specific and evolves with sectoral output according to

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

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). This specification has two advanatges. First, it allows for a direct matching with empirical papers that study the effects of cumulative output on productivities. Second, the additive formulation admits a balanced growth path.

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 which are assumed equal

[16] Fishing-out effects imply that for a balanced growth path 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.

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.[17]

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, 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 self-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 $\frac{\iota_k(1-\alpha)(1-\varepsilon)}{1-\eta}$.

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). In the absence of such spillovers, 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,

[17] See Subsection D.3 for a detailed discussion.

$\iota_K = 0.21$, and substantial decreasing returns to research, $\eta = 0.41$, further mitigate the reinforcing force of knowledge advances. Finally, the fishing-out effect slows the transition toward corner solutions. These features jointly support the existence of a stable interior BGP.

Learning-by-doing adds another layer of complexity and risk of corner solutions. 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.

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

$$\begin{aligned} H_t &= L_{Ft} + L_{Gt} + L_{Nt}, \\ S &= S_{Ft} + S_{Gt}, \\ C_t &= Y_t - \int_0^1 (x_{Fit} + x_{Git} + x_{Nit})\, di. \end{aligned}$$

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.

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 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. Each period, the government is constrained by an emission limit, Ω_t, and it 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:

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

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

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

$$Gov_t^T \geq \bar{Gov}_{min,t}^T \tag{7}$$

subject to the behavior of firms and households, and feasibility.[18] T_{Rt} summarizes subsidies on research, $T_{Rt} = \tau_{sFt} w_{st} S_{Ft} + \tau_{sGt} w_{st} S_{Gt}$. Constraint (5) is the emission limit. 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. Both labor income taxes, τ_{lt}, and carbon taxes, τ_{Ft}, generate government funds. They have to be at least as high as an exogenous funding requirement, $\bar{Gov}_{min,t}^T$, plus spending on research subsidies, T_{Rt}.[19]

In the counterfactual version of the model, lump-sum taxes are feasible, replicating the standard assumption in the literature. Note that even if equipped with lump-sum taxes, the first-best allocation or social planner allocation, defined in Online Appendix B, is not feasible due to learning-by-doing which additionally necessitates sector-specific production subsidies.

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

[19] I do not model government consumption explicitly which would not enter household utility. This ensures that government consumption generates household utility. If resources were wasted, instead, government consumption would affect households' labor supply decision.

3 Theory

This section derives analytical results that guide the interpretation of the quantitative findings. I first highlight the dual role of carbon taxes and their significance in distortionary fiscal settings. Second, I characterize the *expertise effect*, through which fiscal distortions can persistently alter the optimal allocation of researchers.

Equilibrium allocation of researchers. In equilibrium, the demand for researchers equates marginal cost and marginal private benefit. Assuming equal capital shares across sectors, $\alpha_G = \alpha_F = \alpha$, the ratio of green-to-fossil researchers can be written as

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

For the derivation see Subsection A.1. This expression highlights the key determinants of the direction of innovation: relative subsidies, relative prices, market size, knowledge spillovers, and the marginal productivity gains from new technologies.[20]

Substituting equilibrium conditions and the productivity link between knowledge and expertise yields the following equilibrium condition in terms of the carbon tax, research subsidies, expertise, and knowledge:

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta} = (1+\widehat{\tau}_{Ft})_e^{\varepsilon} \frac{1-\tau_{sFt}}{1-\tau_{sGt}} \left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)} \left(\frac{Q_{Gt}}{Q_{Ft}}\right)^{\iota_L(1-\alpha)(\varepsilon_e-1)} \left(\frac{K_{Gt}}{K_{Ft}}\right)^{\iota_L(1-\alpha)(\varepsilon_e-1)-1}. \tag{9}$$

Equation (9) will be the workhorse to later isolate how carbon pricing, research subsidies, and expertise determine the (mis-)allocation of researchers under fiscal distortions.

The effect of a carbon tax and the expertise effect Carbon taxation affects the allocation of researchers directly and indirectly. It directly raises the relative price

[20] The marginal effect on labor-augmenting productivity, equals unity in models assuming a one-to-one linkage between research and productivity Acemoglu et al. (2012); Fried (2018). Assuming a non-linear relationship between productivity and technology, an additional term shows up.

of fossil goods, shifting demand toward green goods. When green and fossil energy are gross substitutes ($\varepsilon_e > 1$), this price effect reallocates demand with elasticity above unity, increasing the market size of the green sector and thus the return to green innovation.

Learning-by-doing introduces an additional channel through which carbon taxes more indirectly affect the direction of innovation. The direction depends on the relative strength of price and market-size effects. More precisely, as green output rises, green expertise Q_G accumulates faster than fossil expertise Q_F. Higher relative expertise raises labor productivity with elasticity ι_L, which increases the market share of green production, feeding back into research incentives. 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 increase in prices and labor supply rises in total by an elasticity of $\iota_L(1-\alpha)(\varepsilon_e - 1)$. All in all, under plausible parameter values, learning-by-doing spurs research in the more experienced sector. 1 summarizes this result.

Proposition 1. *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 $\frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta}$.*

The proof follows directly from eq. (9).

A central implication of Proposition 1 is that any force that changes the *composition* of production—and thus the relative accumulation of expertise—feeds back into the equilibrium incentives that govern the allocation of research. Learning-by-doing translates this relation into a dynamic mechanism. Relative to a first-best implementation of emission targets, fiscal distortions imply an implementation at a smaller scale and lower ratio of green-to-fossil output. As a result, the equilibrium expertise ratio $\frac{Q_{Gt}}{Q_{Ft}}$ reduces, shifting the optimal allocation of researchers persistently to more fossil research.

Another aspect worth emphasizing is the direct effect of carbon taxes on the direction of innovation. At first glance, this mechanism suggests a double dividend of carbon taxation as it not only lowers the externality directly but also fosters green technology

advances. However, this double benefit hinges on the alignment between the optimal transition of expertise and the optimal allocation of research.

In general, this alignment breaks down. In particular, cross-sectoral knowledge spillovers mean that fossil innovation remains valuable in a green future as it stimulates green research. Therefore, as we will see in the quantitative results, the planner prefers to maintain more fossil research than implied by the carbon taxes necessary to implement the emission target. Subsidies on fossil research become attractive as an instrument to counteract this force, but their use is limited by fiscal distortions.

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 balanced growth path (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 achieve the emission target mean the economy is no longer on a BGP with constant ratios. I first discuss the data sources for the emission 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-targeted model outcomes to the data.

4.1 Data

Emission target. I focus on CO_2 emissions and omit other greenhouse gases as they are the largest contributor to warming (IPCC, 2022, 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.[21]

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. 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 PATSTAT database. I filter technologies into green and fossil-relates based on classifications developed in a joint study by the International Energy Agency and the EPO.[22] Fossil energy patents relate to the supply, exploration, processing, transport, and distribution of fossil fuels.[23] 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 2023 for granted patents.[24] Knowledge stocks result as the cumulative patent counts accounting for depreciation. Subsection D.1 details the data collection and the derivation of knowledge stocks.

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

[22] 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.

[23] These are technology 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.

[24] I focus on granted patents to ensure a certain quality of patents.

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%.[25] 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.[26] 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 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 D.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 US are respected.[27] I use

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

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

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

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.[28] 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 targeted grid-search procedure, ensuring the model reproduces empirical patterns in knowledge stocks.

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 scientist supply, 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

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

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 $\tau_{F0} = 0$, reflecting the absence of a federal carbon price in the U.S. in 2015. I also set subsidies on fossil research, $\tau_{sF0} = 0$, 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 D.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, Af/Ag, 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.

As further BGP targets, I employ the average expenditure share on energy relative to GDP equal to 6% from the US Energy Information Administration (EIA, 2023, Table 1.7) which informs[29] the weight on energy in production. The resulting weight on energy is $\delta_y = 0.33$.[30] 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

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

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

tables of the Bureau of Labor Statistics (BLS).[31] 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) used a slightly higher value of 0.91.

Equipped with the parameter values that result from the BGP macthing, I use base-year information to calibrate the labor tac and the emission intensity of fossil energy $\{\tau_{l0}, \omega\}$. To match τ_{l0}, I use observed government expenditures minus new debt relative to GDP in 2015-2019,[32] rationalizing $\tau_{l0} = 0.26$; a value consistent with the literature (Barrage, 2020). To find the parameter relating CO_2 emissions[33] from energy use in the model, ω, I match gross emissions in the base period according to $Emissions = \omega F_0$. The resulting emission intensity of fossil-fuel production is $\omega = 191$. Initial productivity levels $\{A_{F-1}, A_{N-1}, A_{G-1}\}$ follow from normalizing total output in 2015-19, Y_0, to unity and take it as the numeraire: $p_y = 1$ in all periods, and the base-year ratio of fossil-to-green energy informs base-year productivity ratios (EIA, 2023, Table 1.1) as explained in the data section. This concludes the calibration of the non-research block of the model.

Growth parameters The remaining parameters govern technological change: $\{\gamma, \eta, \phi, \nu, \tau_{sG0}, \kappa, K_{G0}, K_{F0}, Q_{G0}, Q_{F0}\}$, 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, τ_{sG0}. For each candidate pair $\{\eta, \tau_{sG0}\}$, 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 stock 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

[31] Retrieved from https://www.bls.gov/green/home.htm, 06 September 2023. The data only contains information for the late 2000s. Ideally, I would use a time series covering the 1990-2010 period.

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

[33] The data comes from EPA (2022) and is described in the data section.

minimization routine to find the allocation. The constraint is to choose a productivity scaling factor κ, which converts units of K and Q 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 $\{\eta, \tau_{sG0}\}$, 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. 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 US 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 $\phi^{AAHK} = 0.3124$.[34] 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

[34]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.

in the range of values in the literature.[35] 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 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 $\nu = 0.16$.

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		0.44
Scientists' productivity γ		1.12
Fishing-out effect ν		0.16
Initial knowledge stock ($K_{F2010-14}$, $K_{G2010-14}$)	knowledge stock in 2010-2014	(1.25, 1.00)
Initial know-how ($Q_{F2010-14}, Q_{G2010-14}$)	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 final good 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)
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

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

Model fit Figure 1 contrasts model outcomes to the data for non-targeted in the upper and targeted time series in the lower panel. Figure 1a depicts gross emissions. These variables are not targeted in the calibration. The model captures dynamics of emissions fairly well until the early 2000. Afterwards, policy interventions 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 7GtCO_2-equivalent in 2010, which exceeds observed emissions by 27%.

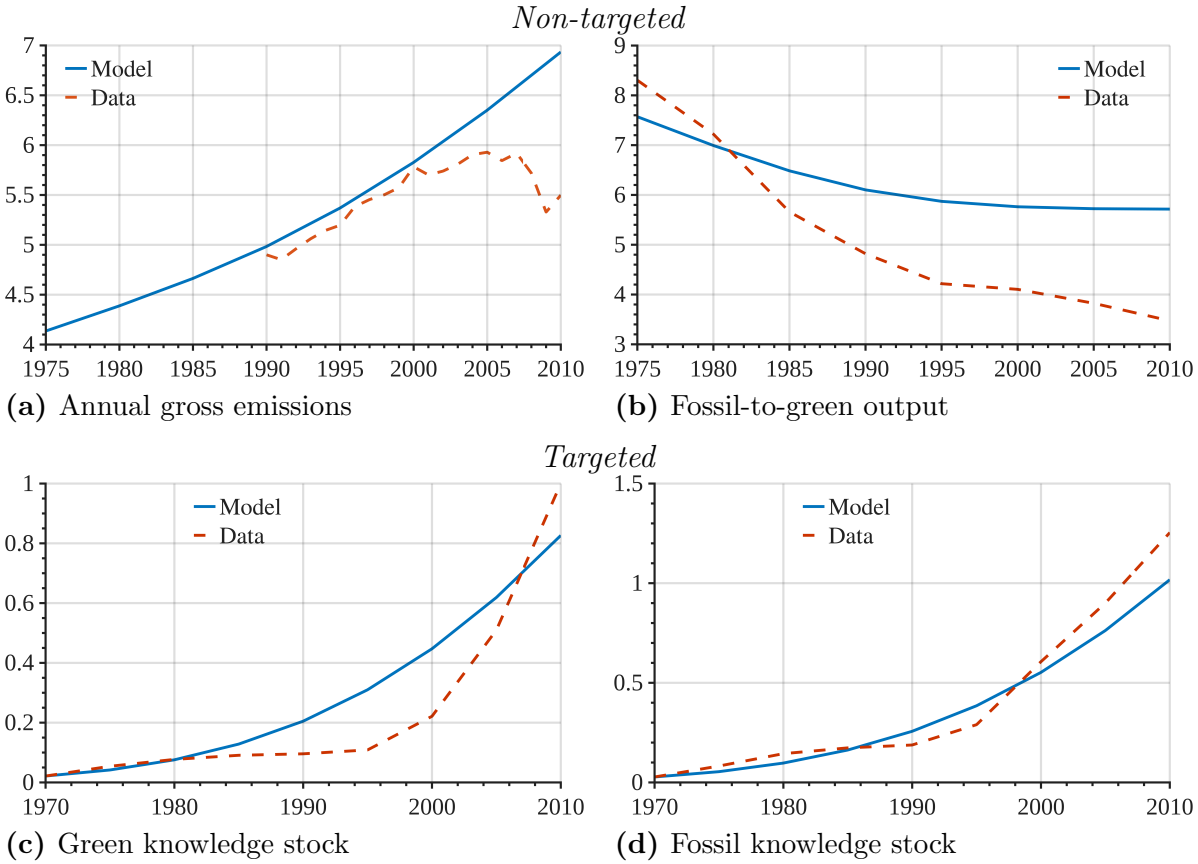
Figure 1b presents the comparison of the model and data-implied ratio of fossil-to-green output. The model replicates the steady decline in the ratio of fossil-to-green energy, though the decline is somewhat muted relative to the data. The non-modeled oil-price shocks of the 1970s and regulatory changes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

Finally, the lower panel of Figure 1 displays the evolution of knowledge stocks in the green (Figure 1c) and fossil sector (Figure 1d), both of which are targeted in the calibration exercise. The model reproduces the convex growth of knowledge but falls short of capturing the acceleration around 2010. Taken together, these figures provide reassurance that the model captures both the broader emission dynamics and the underlying trends in knowledge accumulation reasonably well.

5 Results

This section presents the main quantitative results. I first describe the optimal environmental policy mix in the baseline regime and its implications for output and innovation along the transition path (Subsection 5.1). I then analyze the mechanisms through which fiscal distortions and learning-by-doing shape the optimal fiscal mix (Subsection 5.2). Next, I examine the implications for the allocation of research and quantify the resulting misallocation under fiscal constraints (Subsection 5.3). Finally, I assess the welfare costs of fiscal distortions (Subsection 5.4)and discuss the sensitivity of the main results (Subsection 5.5).

Figure 1: Model Evaluation.



Notes: Panels (a)–(b) report non-targeted variables for the model and from the data. 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. Panels (c)–(d) compare the model-implied evolution of the green and fossil knowledge stock to the data.

5.1 The optimal policy mix

Figure 2 summarizes the main quantitative results. Panel 2b reports the optimal carbon tax in levels, while Panel 2a shows the carbon tax relative to the Pigouvian benchmark. Panels 2c and 2d report, respectively, the optimal labor income tax and net research subsidy spending, each expressed as deviations from their optimal levels in an otherwise identical economy without environmental externalities. Panels 2e and 2f display the implied allocation outcomes, measured 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 Panel 2b, the carbon tax rises from an initial level of \$60 to about \$480 by 2070 (2022 prices). Relative to the

Pigouvian benchmark—defined as the marginal social cost of carbon[36]—the carbon tax exceeds the benchmark by about 3.5% initially (Panel 2a). The wedge narrows to roughly 2.5% by 2030 and remains positive thereafter. Hence, throughout the transition, the optimal carbon tax remains above the Pigouvian level despite the presence of distortionary taxation.

The labor income tax declines relative to its non-target optimum (Panel 2c). Initially, the reduction amounts to 4.2 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 decline while government spending requirements grow with output.

Net expenditures on research subsidies are substantially reduced in the early phases of the transition. Relative to the non-target benchmark, net R&D subsidy spending falls by 69% initially (Panel 2d). Only around 2040–2045 do research subsidies exceed their non-target level, and in the long run they are about 60% higher than in the non-target optimum. Hence, while reliance on carbon taxation is strongest early in the transition, research subsidies gain importance at later stages. These dynamics suggest a shift over time from policies that primarily steer production toward green technologies to policies that increasingly guide the allocation of research.

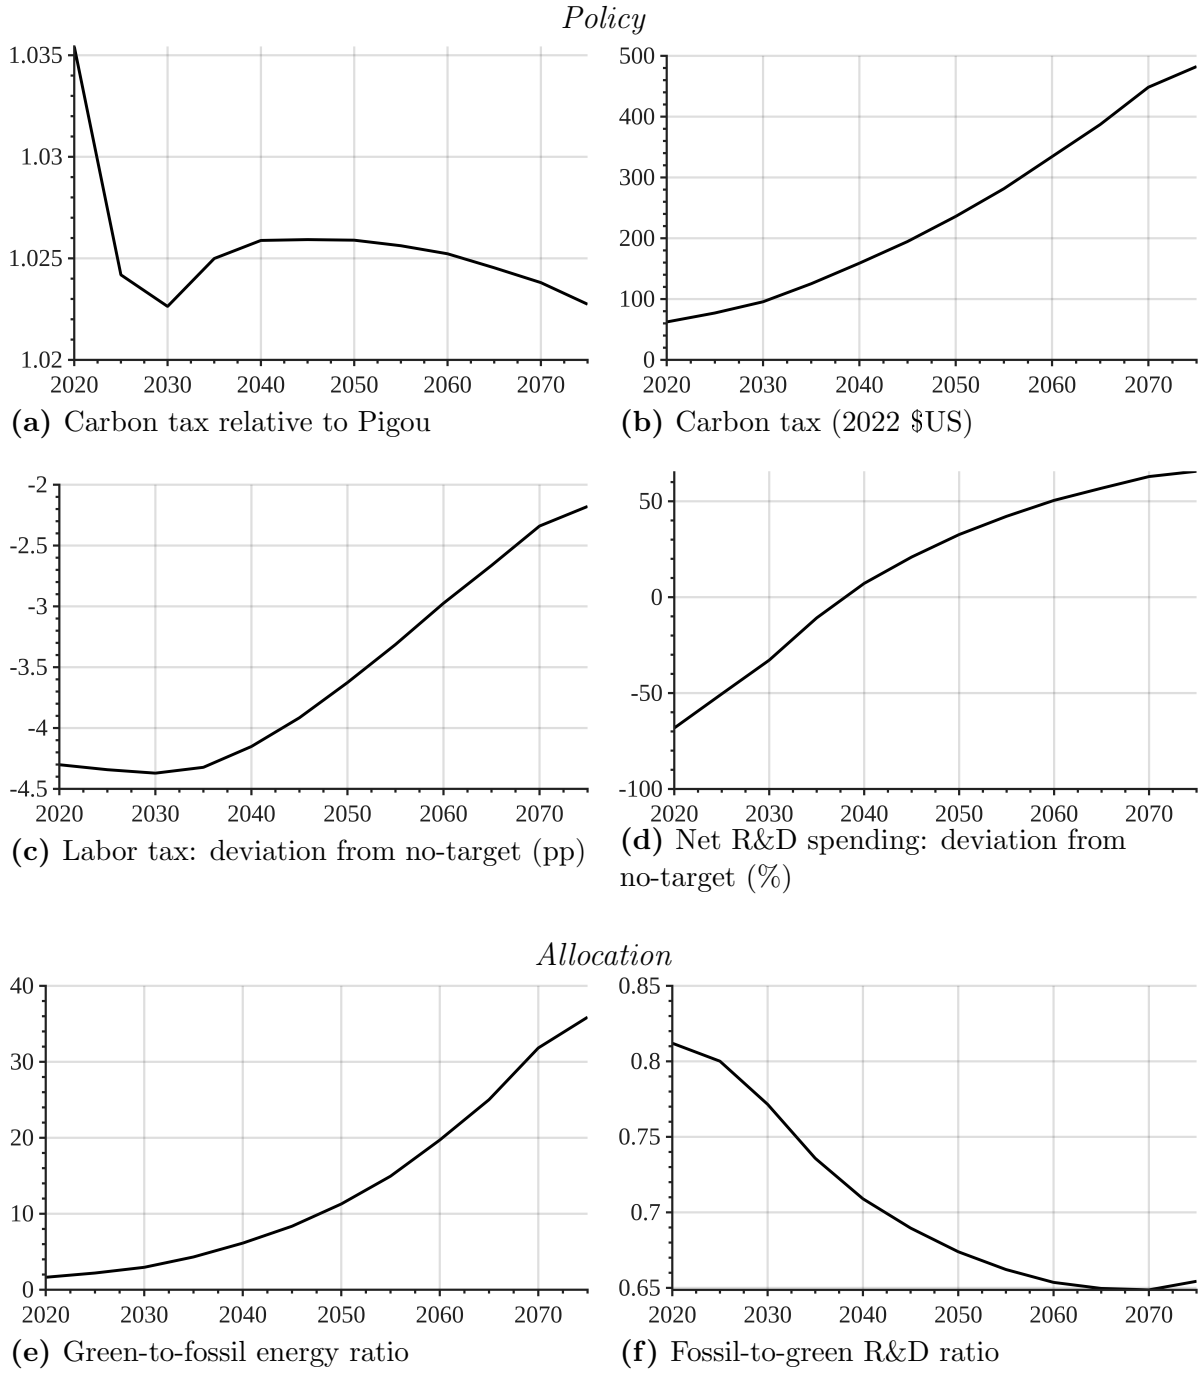
Under this policy mix, the emissions target is exactly met. The green-to-fossil energy ratio rises monotonically over time (Panel 2e), while the fossil-to-green R&D ratio declines from about 0.8 initially to roughly 0.65 in the long run (Panel 2f).

5.2 Fiscal Distortions and Learning-by-doing: The optimal fiscal mix

This section clarifies how fiscal distortions shape the optimal policy mix and how learning-by-doing modifies these second-best effects. I proceed in three steps. First, I analyze a benchmark economy with lump-sum taxation and endogenous learning-by-doing, which

[36]In the numerical exercise, the Pigouvian benchmark is given by the shadow value of the emissions constraint in the Ramsey planner's problem.

Figure 2: Optimal policy mix.



(a) Carbon tax relative to Pigou

(b) Carbon tax (2022 $US)

(c) Labor tax: deviation from no-target (pp)

(d) Net R&D spending: deviation from no-target (%)

(e) Green-to-fossil energy ratio

(f) Fossil-to-green R&D ratio

Notes: The figure reports the jointly optimal time paths for policy instruments and allocations under the baseline calibration. Panels 2a–2b show the carbon tax relative to the Pigouvian benchmark and in levels. Panels 2c–2d report the labor income tax and net research subsidy spending, each expressed as deviations from their optimal levels in an otherwise identical economy without environmental externalities. Panels 2e–2f display the implied allocation outcomes in terms of energy production and research effort.

yields the policy mix I refer to as *first best*.[37] Second, I introduce distortionary taxation

[37]The policy is first-best in the sense that it refers to a setting without fiscal distortions. However, due to the absence of targeted technology subsidies, the resulting allocation is not fully efficient.

while holding learning fixed at its optimal path under lump-sum finance, so that the planner faces fiscal distortions but learning dynamics are exogenous. Third, I allow learning-by-doing to adjust endogenously under distortionary taxation, which yields the baseline model. While comparing the lump-sum benchmark to the baseline model is informative about the overall importance of fiscal distortions, comparing the latter two environments isolates how endogenous learning affects second-best policy design.

Figure 3 provides a visual summary. Panels 3a–3b report the carbon tax relative to the Pigouvian benchmark and in levels. Panels 3c–3d report the labor income tax and net research subsidy spending, each expressed as deviations from their non-target optima. Panels 3e–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 fiscal-distortion regime with learning held fixed at its benchmark path (orange dash-dotted), and the baseline with fiscal distortions and endogenous learning-by-doing (black solid).

Fiscal distortions: A less distortionary fiscal mix. I first compare the regime with fiscal distortions and fixed learning to the lump-sum benchmark to isolate the role of fiscal distortions. In the presence of distortionary taxation, the carbon tax required to meet the emissions target is lower in levels (Panel 3b) and lower relative to the Pigouvian benchmark (Panel 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. The effect on labor supply arises through general-equilibrium mechanisms: by raising production costs and shifting the input mix, carbon pricing lowers the after-tax wage and thus depresses labor supply. 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.

In line with the weak double-dividend logic, carbon tax revenues are 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 the transition (Panel 3c), with the reduction being largest early on when carbon tax revenues are highest.

Fiscal distortions also alter the scale and composition at which the emissions target is implemented. With distortionary taxation, the target is achieved at a smaller output scale and with a higher fossil share than under lump-sum finance, reflected in a lower green-to-fossil energy ratio (Panel 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 matters for the desired allocation of research.

Learning-by-doing: A more distortionary fiscal mix. I next discuss how learning-by-doing modifies second-best policy design by comparing the baseline (black solid) to the fiscal-distortion regime with fixed learning (orange dash-dotted). Most notably, allowing learning to adjust endogenously raises the optimal carbon tax relative to the fixed-learning case and reverses the Pigouvian wedge: in the baseline, the carbon tax exceeds the Pigouvian benchmark (Panel 3a). In the absence of targeted instruments to directly subsidize learning, carbon pricing becomes an instrument to shift production—and thus learning—toward the green sector. This motive is quantitatively strongest early in the transition, when the value of building green expertise is highest.

This result is noteworthy because the public-economics literature emphasizes that learning-by-doing can raise the distortionary cost of fiscal instruments by affecting not only current labor supply but also future productivity (Peterman, 2016). Despite this force, the planner is willing to tolerate a more distortionary fiscal mix—higher reliance on carbon pricing and correspondingly lower labor taxation—because what matters during the transition is not only the level of economic activity, but where expertise is accumulated.

When lump-sum taxes are feasible, the planner can additionally use lump-sum instruments to support labor supply; in that benchmark, the labor income tax can turn negative in levels (not shown), consistent with the large reductions relative to the non-target optimum

shown in Panel 3c.[38] In the distortionary regime, such a labor subsidy is infeasible. Nevertheless, in the baseline, higher carbon tax revenues allow for a larger reduction in the labor income tax relative to its non-target optimum (Panel 3c).

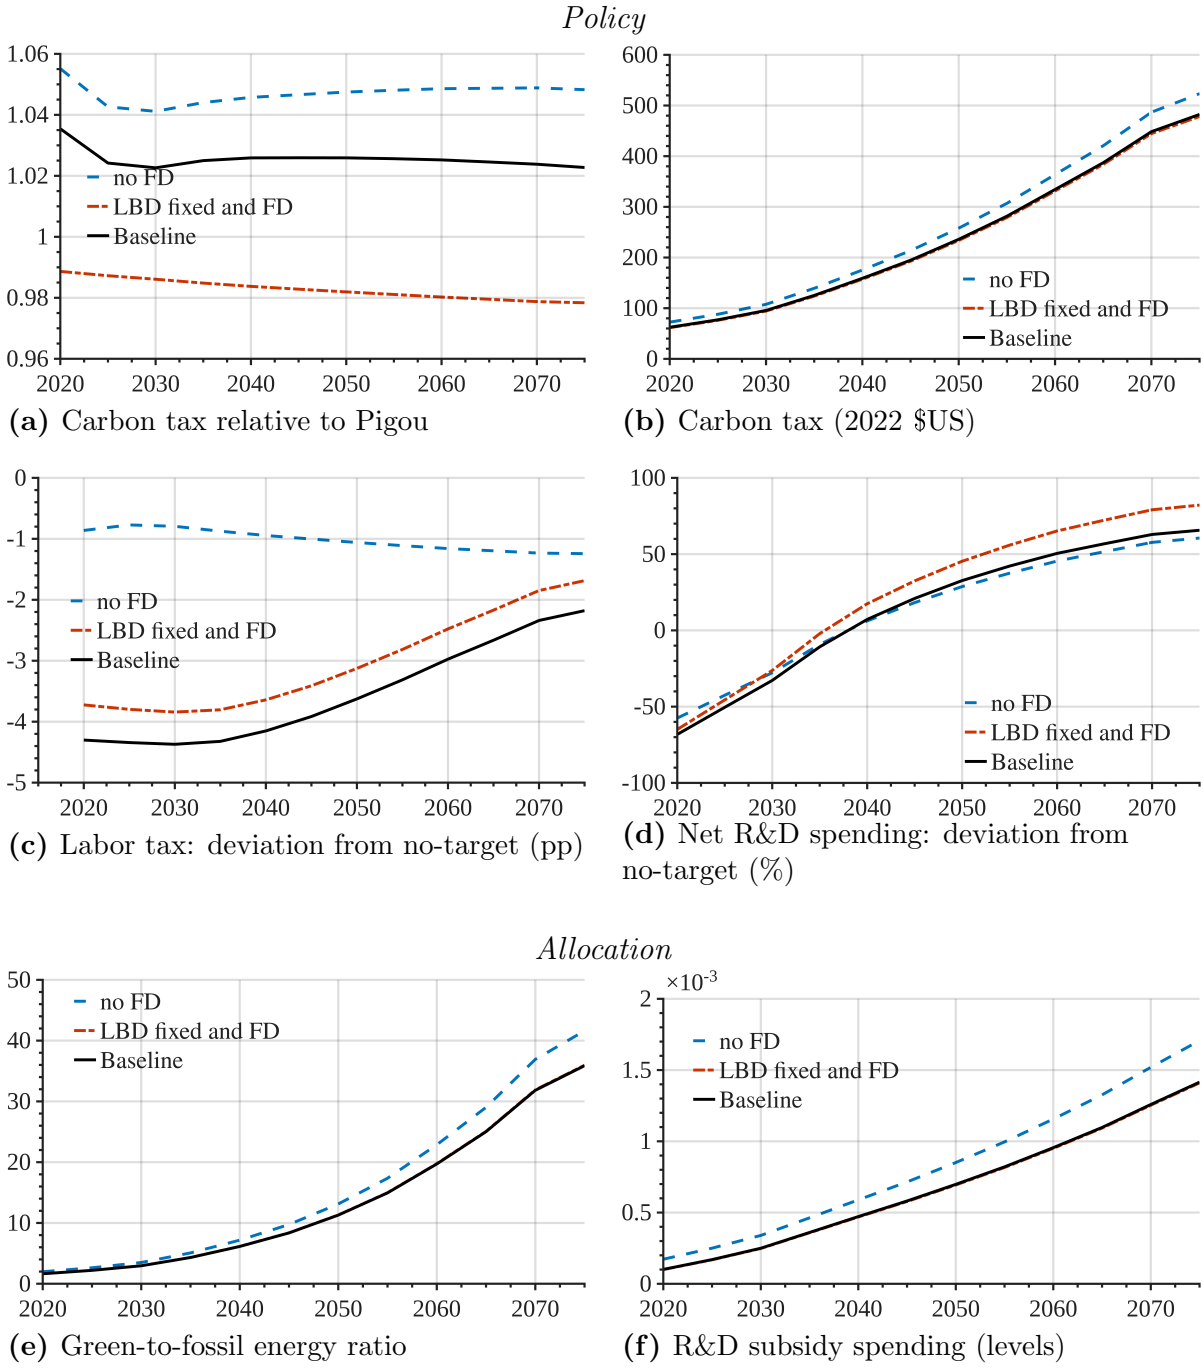
Expenditures on research subsidies. I close this section by turning to research subsidies. For consistency with Subsection 5.1, Panel 3d reports net research subsidy spending as a deviation from the non-target optimum. While these deviation paths are similar across regimes, they mask a lower level of expenditures under fiscal distortions (Panel 3f). At this stage, it is unclear whether lower spending primarily reflects a fiscal motive—reducing reliance on distortionary taxation—or a change in the desired allocation of research induced by the altered scale and composition of production under fiscal distortions. I address this distinction explicitly in Subsection 5.3. I also examine more closely the role of learning-by-doing: on the one hand, learning-by-doing raises the efficiency costs of reduced labor supply, which may call for lower research subsidy expenditures; on the other hand, a misalignment between the optimal trajectories of learning and innovation could increase the value of research subsidies. These opposing forces may partly offset each other in Figure 3.

5.3 Misallocation of R&D and learning-by-doing

I now examine how fiscal constraints distort the allocation of research effort between green and fossil technologies. In the preceding analysis of the optimal policy mix in Subsection 5.2, the planner scales back research subsidies in response to fiscal distortions, while endogenizing learning-by-doing has little effect on subsidy expenditures in levels. As this section shows, fiscal distortions reduce the need to finance research subsidies but nevertheless induce a misallocation of research. Moreover, when learning-by-doing is endogenous, the planner accepts a more severe misallocation in order to limit additional efficiency costs, consistent with the innovation–learning dilemma emphasized throughout

[38] Comparing these results to the first-best allocation shows that even without explicit instruments to target sectoral learning-by-doing, such as technology subsidies, the combination of an above-Pigouvian carbon tax and a labor subsidy closely approximates the first-best outcome where learning spillovers can be fully internalized.

Figure 3: Mechanism: fiscal distortions and learning-by-doing.



(a) Carbon tax relative to Pigou

(b) Carbon tax (2022 $US)

(c) Labor tax: deviation from no-target (pp)

(d) Net R&D spending: deviation from no-target (%)

(e) Green-to-fossil energy ratio

(f) R&D subsidy spending (levels)

Notes: The figure compares three environments: (i) a lump-sum benchmark with endogenous learning-by-doing, (ii) a regime with fiscal distortions and learning fixed at the benchmark path, and (iii) the baseline with fiscal distortions and endogenous learning-by-doing. All deviation panels are expressed relative to the corresponding no-target optimum in the same environment. Line styles: lump-sum benchmark (blue dashed), fixed-learning fiscal regime (orange dash-dotted), baseline (black solid).

the paper.

Measuring misallocation of research. To gauge the extent of misallocation, 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 of research misallocation. 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 at a lower overall output level, most likely raising the social value of fossil research. A direct comparison between first-best and second-best outcomes therefore conflates changes in the planner's target research mix with limitations imposed by fiscal constraints.

To disentangle these effects, I conduct a diagnostic 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 energy production—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 *desired* allocation of research. Misallocation arising from fiscal distortions is then given by the difference between the fully distortionary regime and the desired allocation.

Funding requirements and misallocation. This experiment is informative about the role of fiscal constraints for research spending. Figure 4b shows R&D subsidy expenditures under fiscal distortions, under full lump-sum taxation, and under lump-sum–financed subsidies. Most of the reduction in research spending under fiscal distortions is explained by a lower requirement to finance subsidies, rather than by an inability to use them: compare the non-distortionary case (blue solid) to the lump-sum–financed subsidy benchmark (orange dash-dotted).

Despite this reduced need for funding, the planner nevertheless accepts a misallocation of research. Figure 4a displays the allocation of researchers as the ratio of fossil-to-green scientists across regimes. The desired allocation under fiscal distortions features a higher

fossil share than in the non-distortionary setting, reflecting the higher fossil share of energy production. When equipped with lump-sum–financed subsidies, the planner would increase the fossil share of research relative to the fully distortionary regime, indicating that the Ramsey planner forgoes a more productive allocation of research in order to reduce the fiscal burden.

Figure 4c quantifies this 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. The misallocation ranges between 3.3% and 4.2% during the transition.

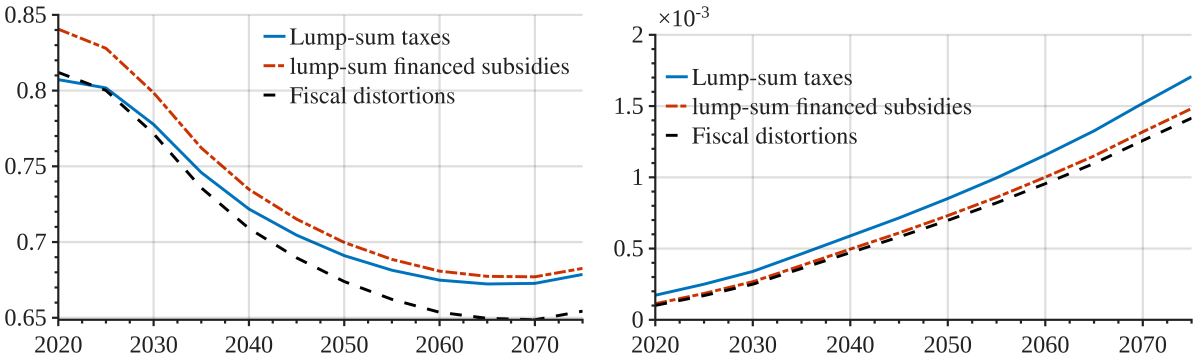
Complementarity of carbon taxes and research subsidies. The direction of misallocation is shaped by the dual role of research. Fossil R&D not only raises fossil productivity but also generates knowledge that spills over into green innovation. These cross-sectoral spillovers make it socially desirable to maintain fossil research during the transition—not for its direct technological contribution, but because it enriches the knowledge base that supports green innovation.[39] This maintained value of fossil research conflicts with the urgency to shift production toward green energy sources. At the same time, market forces redirect 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, however, the carbon tax required to meet emissions targets is lower, which already dampens the shift of innovation toward the green sector and reduces the marginal usefulness of research subsidies.

Learning-by-doing and the misallocation of research. The same mechanism is crucial for understanding how learning-by-doing affects the misallocation of research. As shown in Subsection 5.2, endogenous learning-by-doing raises the value of directing production toward green technologies and motivates the use of higher carbon taxes. While

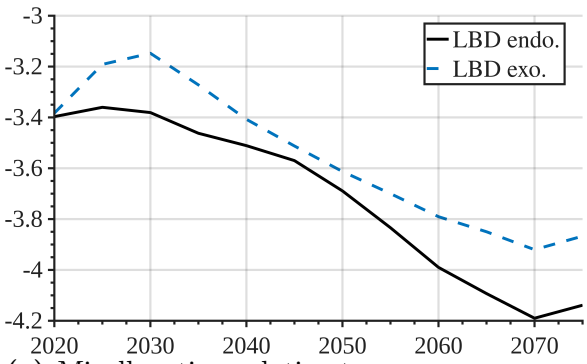
[39] Barbieri et al. (2023) provide empirical evidence for positive spillovers from fossil to green innovation.

Figure 4: Misallocation of research under fiscal distortions and learning-by-doing.

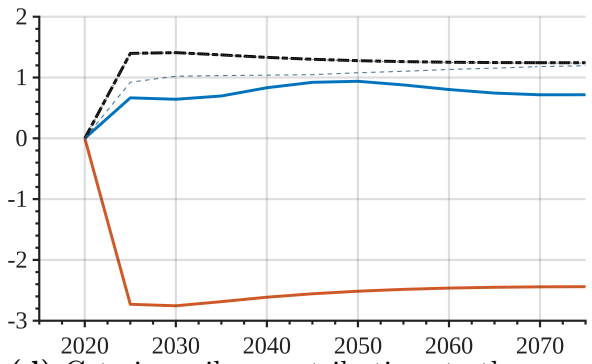


(a) Fossil-to-green researchers

(b) R&D subsidy spending (levels)



(c) Misallocation relative to lump-sum–financed subsidies



(d) Ceteris paribus contributions to the research ratio

Notes: Panel (a) shows the ratio of fossil-to-green researchers under different fiscal regimes in the baseline model. For the same regimes, Panel (b) reports R&D subsidy expenditures in levels. Panel (c) quantifies the misallocation of research induced by fiscal distortions, measured as the percentage deviation of the fossil-to-green research ratio in the fully distortionary regime from a benchmark with lump-sum financed research subsidies. Panel (d) decomposes the difference in the logged 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 based on eq. (8).

this policy response supports the accumulation of green expertise, it also intensifies the conflict with the objective of maintaining fossil research for its spillover benefits. Learning-by-doing therefore alters not only the level of optimal policy instruments, but also the trade-off between directing innovation and preserving a diversified knowledge base. Whether this mechanism amplifies or mitigates misallocation is ultimately an empirical question. Comparing the baseline misallocation to the same metric arising in the model where learning is exogenous, reveals that learning-by-doing amplifies the

misallocation.[40] As shown by the dashed-blue line in Figure 4c, misallocation declines by between 8.5% and 12% when learning-by-doing is exogenous.

To understand the underlying mechanisms, I return to the general equilibrium condition that determines the allocation of research effort between fossil and green technologies eq. (8)). Taking logs of the inverse condition yields an additive expression for the log research ratio, $\log(s_f/s_g)$. 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 fully 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% and 3%. Second, the expertise effect works in the opposite direction: because fiscal distortions imply a less green production mix raising the share of fossil research by slightly more than 1%, ceteris paribus. Third, research subsidies adjust to counteract the higher carbon tax, but at less than 1%.

The thin lines in Figure 4d show the same decomposition under lump-sum–financed subsidies, corresponding to changes in the *desired* 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.

In sum, the value of fossil knowledge complicates the task of directing learning-by-doing toward green technologies. Policies that foster green learning reinforce private incentives to concentrate research in the green sector. Because research subsidies are fiscally costly, this concentration cannot be fully undone. The additional value of research subsidies when learning is endogenous helps explain why research subsidy expenditures

[40] In line with the previous sections, learning-by-doing is fixed at the optimal level in the first-best regime.

do not decline more to mute fiscal distortions in Figure 3f.

5.4 Welfare analysis: Cost of fiscal distortions

This section aggregates the mechanisms discussed above into a welfare metric. The welfare analysis serves two purposes: to quantify the aggregate cost of fiscal distortions and to assess whether the stronger misallocation of research under learning-by-doing is also more costly in welfare terms. I quantify the welfare cost of financing the green transition through distortionary taxation—and how this cost is affected by learning-by-doing—using the consumption equivalent variation (CEV). The CEV measures the fraction of per-period consumption that the representative household would be willing to permanently forgo to avoid fiscal distortions.

The results are summarized in Table 2. When an emission target is imposed, the welfare loss from fiscal distortions amounts to 1.3% of per-period consumption when learning-by-doing is exogenous and rises to 1.5% when learning-by-doing is endogenous, corresponding to a 15% amplification. Without an emission target, the amplification is more modest, at around 8%. These patterns highlight two distinct channels through which learning-by-doing raises the welfare cost of fiscal distortions.

First, even in the absence of an emission target, fiscal distortions are more costly when learning-by-doing is endogenous because they interfere with dynamic productivity gains. Distortionary taxation depresses labor supply and production, which slows the accumulation of experience and raises long-run efficiency costs. This channel explains why welfare losses from fiscal distortions are higher under endogenous learning even in the no-target environment.

Second, during the green transition, endogenous learning-by-doing strengthens the case for using carbon taxes to direct production—and thus learning—toward green technologies. As shown in Subsection 5.2, this reliance on carbon pricing becomes more pronounced when learning-by-doing is endogenous. Because carbon taxes are more distortionary than labor taxes, their more aggressive use further amplifies the welfare cost of fiscal

distortions when an emission target is in place.[41]

Finally, Table 2 also reports the welfare cost associated with the misallocation of researchers. Correcting the misallocation yields welfare gains that are 11% larger in the model with learning-by-doing than in the model with exogenous learning. This finding reinforces the results from Subsection 5.3: learning-by-doing calls for higher carbon taxes and a muted response of research subsidies, which amplifies the misallocation of research effort in levels, making lump-sum financed research subsidies more valuable.

Table 2: Welfare loss from fiscal distortions as percent of per-period consumption (CEV)

	Exo. LBD	**With LBD**	**LBD/Exo. LBD**
With Target	1.30%	1.50%	15%
No Target	1.71%	1.85%	8%
R&D Misallocation	0.0054%	0.0059%	11%

Notes: The first two rows are measured relative to the non-distortionary first-best allocation. *R&D Misallocation* refers to the gap between the fiscal distortion regime and a regime with lump-sum-financed research subsidies.

5.5 Robustness

To assess the stability of the main results, I conduct a series of targeted robustness checks. Each exercise varies a single parameter while holding the rest of the calibration fixed. I focus on five key dimensions: (i) the elasticity of learning-by-doing, governed by ι_L; (ii) the strength of cross-sectoral knowledge spillovers, ϕ; (iii) the returns to research, η; (iv) the labor share in the green sector, α_g; and (v) policy targets, including the level of exogenous fiscal pressure, $\bar{G}^T_{\min,t}$, and the emissions constraint. For each case, I compare the carbon tax relative to the Pigouvian benchmark and the misallocation of research due to fiscal distortions, including its amplification through the interaction with learning-by-doing.

The results confirm that learning-by-doing consistently pushes the carbon tax above

[41] I also decompose the CEV loss in the baseline model with an emission target into three components: (i) fiscal distortions when learning-by-doing is fixed, (ii) the additional contribution from endogenizing learning-by-doing, and (iii) an interaction term. The decomposition shows that fiscal distortions under exogenous learning account for 87% of the total welfare loss, while endogenizing learning-by-doing contributes an additional 13%, consistent with the comparison between the endogenous- and exogenous-learning environments.

the social cost of carbon. This wedge is most pronounced in the short run and declines over time, remaining positive in the long run. A lower capital share (i.e., higher labor share) in the green sector strengthens this mechanism, as workers experienced in green technologies become more important for green energy production.

Results are surprisingly insensitive to the strength of cross-sectoral spillovers. However, eliminating spillovers altogether leads to a further increase in the carbon-to-Pigouvian ratio, consistent with the idea that persistent fossil knowledge dampens the planner's incentive to accelerate green learning to avoid pulling research to the green sector. A lower η (higher duplication risk) allows for more aggressive rise in learning-by-doing in the green sector, as researcher allocation becomes less sensitive to expertise gaps.

Higher fiscal pressure reduces the carbon tax wedge, consistent with the view that directing learning through carbon pricing is fiscally costly. Imposing a net emissions target aligned with a 1.5°C pathway raises the carbon tax above the Pigouvian benchmark in the short run, while narrowing the wedge in the long run. This pattern highlights the importance of building green learning early in the transition as climate targets tighten.

Turning to research misallocation, the robustness checks confirm that misallocation ranges from 3–5% and is amplified by learning-by-doing by 3–22%. One notable exception arises when cross-sectoral spillovers are shut down. Then learning-by-doing mitigates misallocation as it directs research to the green sector.

6 Conclusion

The transition to net-zero emissions demands a reorganization of production and innovation. While motivating green learning and innovation is central to the transition to net-zero emissions, much less is known about how these objectives should be targeted when climate policy is constrained by distortionary taxation.

This paper studies optimal climate policy in a dynamic general-equilibrium model with learning-by-doing, directed technical change, and distortionary taxation. The analysis highlights an innovation–learning dilemma of fiscal policy 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, fiscal distortions both shift the socially desirable allocation of research and induce a persistent misallocation as research subsidies are fiscally costly. Third, a welfare decomposition shows that learning-by-doing amplifies the welfare costs of fiscal distortions and increases the losses associated with research misallocation.

References

Acemoglu, D. (2002). Directed Technical Change. Review of Economic Studies, 69:781–809.

Acemoglu, D., Aghion, P., Barrage, L., Hémous, D., and Zurich, L. B. E. (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. Technical report.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, 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.

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).

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.

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. Working Paper 33378, National Bureau of Economic Research.

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

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.

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, 57:369–391.

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.

Dasgupta, P. (2021). The Economics of Biodiversity: The Dasgupta Review. Technical report, London: HM Treasury.

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

Douenne, T., Hummel, J. A., and Pedroni, M. (2022). Optimal Fiscal Policy in a Climate-Economy Model with Heterogeneous Households.

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: 2026-02-09.

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.

Hart, R. (2019). To Everything There Is a Season: Carbon Pricing, Research Subsidies, and the Transition to Fossil-Free Energy. The Association of Environmental and Resource Economists, 6:135–175.

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

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.

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

IPCC (2022). *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [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.

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.

Loebbing, J. (2019). National Tax Association Redistributive Income Taxation with Directed Technical Change. 112:1–97.

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. (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, pages 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: 2026-02-09.

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.

World Commission on Environment and Development (1987). *Our Common Future*. New York: Oxford University Press.

Young, A. (1993). Invention and bounded learning by doing. Journal of Political Economy.

Zwaan, B. C. C. V. D., Gerlagh, R., Klaassen, G., and Schrattenholzer, L. (2002). Endogenous technological change in climate change modelling. Energy Economics, 24:1–19.

A Model equations and simplifications

This section details the budget of the government and the household before I spell out the equilibrium equations.

Government: Budgets

$$\begin{aligned}
\text{Labor Income Taxation} \quad & \tau_{lt} w_t H_t = T_{Lt} \\
\text{Environmental Policy} \quad & \tau_{Ft} F_t = T_{Ft} \\
\text{Research Subsidization} \quad & T_{Rt} = w_{st} \tau_{sGt} s_{Gt} + w_{st} \tau_{sFt} s_{Ft} \\
\text{Monopoly Correction} \quad & T_{xt} = \int_0^1 (p_{Fit} \zeta_t x_{Fit} + p_{Nit} \zeta_t x_{Nit} + p_{Git} \zeta_t x_{Git}) di \\
\text{Government Budget} \quad & Gov_t^T = T_{Ft} + T_{Lt} - T_{Rt}
\end{aligned}$$

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:

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

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

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

T_{Rt} remains relevant for the planning problem as it shifts government revenue conditions from non-lump-sum taxes.

Household

$$\begin{aligned}
\text{Period utility} \quad & \log(C_t) - \chi \frac{H_t^{1+\sigma}}{1+\sigma} \\
\text{Budget} \quad & C_t = (1 - \tau_{lt}) w_t H_t + w_{st} S + \Pi_t + Gov_t^T - T_{xt}
\end{aligned}$$

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

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

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

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

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

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

Machine producers

Price setting $p_{xFit} = \dfrac{1}{\alpha_F (1+\zeta_F)}$

$p_{xNit} = \dfrac{1}{\alpha_N (1+\zeta_N)}$

$p_{xGit} = \dfrac{1}{\alpha_G (1+\zeta_G)}$

Demand Scientists $w_{st}(1-\tau_{sFt}) =$

$\eta\gamma(1-\alpha_F) \left(K_{Ft-1}^{1-\phi} K_{-Ft-1}^{\phi} \right)^{1-\nu} (s_{Ft})^{\eta-1} \dfrac{p_{Ft} F_t}{K_{Ft}}$

$w_{st}(1-\tau_{sGt}) =$

$$\eta\gamma(1-\alpha_G)\left(K_{Gt-1}^{1-\phi}K_{-Gt-1}^{\phi}\right)^{1-\nu}(s_{Gt})^{\eta-1}\frac{p_{Gt}G_t}{K_{Gt}}$$

Innovation
$$K_{Fit} = K_{Ft-1}(1-\delta_K) + \gamma(s_{Fit})^{\eta}\left(K_{Ft-1}^{1-\phi}K_{-Ft-1}^{\phi}\right)^{1-\nu}$$
$$K_{Git} = K_{Gt-1}(1-\delta_K) + \gamma(s_{Git})^{\eta}\left(K_{Gt-1}^{1-\phi}K_{-Gt-1}^{\phi}\right)^{1-\nu}$$

Government
$$T_{Ft} = \tau_{Ft}F_t$$
$$T_{\text{subst}} = \int_0^1 (p_{xFit}\zeta_{Ft}x_{Fit} + p_{xGit}\zeta_{Gt}x_{Git} + p_{xNit}\zeta_{Nt}x_{Nit})\,di + \int_0^1 (w_{st}\tau_{sFt}s_{Fit} + w_{st}\tau_{sGt}s_{Git})\,di$$

with
$$\zeta_{Jt} = \frac{1-\alpha_J}{\alpha_J} \; for \; J \in \{F,G,N\}$$

Markets
$$L_{Ft} + L_{Nt} + L_{Gt} = H$$
$$s_{Ft} + s_{Gt} = S$$
$$C_t + \int_0^1 (x_{Fit} + x_{Nit} + x_{Git})\,d_i = Y_t$$

A.1 Analytical 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' optimality condition. 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 the demand for machines

$$x_{Jt} = \left(\frac{p_{Jt}\alpha_J}{p_{xJt}}\right)^{\frac{1}{1-\alpha_J}} L_{Jt}A_{Jt}, \tag{10}$$

and labor demand

$$L_{Jt} = (1-\alpha_J)\frac{p_{Jt}}{w_t}J_t. \tag{11}$$

Eq. (10) shows that the elasticity of machine demand with respect to its price is $\frac{1}{1-\alpha_J}$. Substituting (10) 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

$$J_t = (p_{Jt}\alpha_J)^{\frac{\alpha_J}{1-\alpha_J}} L_{Jt}A_{Jt}. \tag{12}$$

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. (10). 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, to indicate 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. (10), demand for machines responds to a marginal rise in s_{Jt} by

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

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 termds, the demand for scientists in equilibrium is given by

$$w_{st}(1-\tau_{sJt})=\frac{(1-\alpha_J)}{\alpha_J}(\alpha_J p_{jt})^{\frac{1}{1-\alpha_J}}L_{Jt}\iota_K\frac{A_{Jt}}{K_{Jt}}\gamma\eta\left(K_{-Jt-1}^{\phi}K_{J_t-1}^{1-\phi}\right)^{1-\nu}S_{Jit}^{\eta-1}. \tag{13}$$

Equilibrium allocation of scientists. The equilibrium allocation of researchers follows by dividing (13) for sector G by sector F and rearranging terms:

$$\left(\frac{S_{Gt}}{S_{Ft}}\right)^{1-\eta}=\frac{\widehat{\alpha_G}}{\widehat{\alpha_F}}\frac{1-\tau_{sFt}}{1-\tau_{sGt}}\left(\frac{p_{Gt}^{\frac{1}{1-\alpha_G}}}{p_{Ft}^{\frac{1}{1-\alpha_F}}}\right)\frac{L_{Gt}}{L_{Ft}}\left(\frac{K_{Gt-1}}{K_{Ft-1}}\right)^{(1-2\phi)(1-\nu)}\left(\frac{A_{Gt}/K_{Gt}}{A_{Ft}/K_{Ft}}\right), \tag{14}$$

where $\widehat{\alpha}_J\coloneqq\alpha_J^{\frac{\alpha_J}{1-\alpha_J}}(1-\alpha_J)$. 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. Labor demand, eq. (11) and production in equilibrium, eq. (12), follows 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 yields:

$$\frac{p_{Ft}^{\frac{1}{1-\alpha_F}}}{p_{Gt}^{\frac{1}{1-\alpha_G}}} = \frac{A_{Gt}\widehat{\alpha}_G}{A_{Ft}\widehat{\alpha}_F}. \tag{15}$$

Note that it follows from this relation of prices and productivity that the price effect and the marginal productivity effect cancel in eq. (14). Expertise affects the allocation of scientists through labor demand. This is why the elasticity of input factors, i.e., labor, determines the stability of an interior BGP.

In general equilibrium, labor demand is determined by demand from energy producers for intermediates. Substituting sectoral output in labor demand (11) 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:

$$\frac{L_{Gt}}{L_{Ft}} = (1+\widehat{\tau}_{Ft})^{\varepsilon_e}\left(\frac{p_{Ft}}{p_{Gt}}\right)^{\varepsilon_e-1}\frac{1-\alpha_G}{1-\alpha_F}. \tag{16}$$

where $\widehat{\tau}_{Ft} \coloneqq \frac{\tau_{Ft}}{p_{Ft}}$. Eq. (16) 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. The carbon tax only affects energy producers' demand leaving intermediate good producers' profits unaffected, ceteris paribus. Hence, independent of whether green and fossil goods are gross substitutes, the carbon tax unambiguously implies a shift towards green production.

Plugging in eq. (15) and eq. (16) in the equilibrium allocation of scientists, eq. (14), yields

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

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 expertise effect It is easy to see that the elasticity of the allocation of researchers to the gap in expertise is given by $\frac{\iota_L(1-\alpha)(\varepsilon_e-1)}{1-\eta}$. Next to the elasticity of substitution of energy inputs, three 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.

B Social planner

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

$$\sum_{t=0}^{\infty} \beta^t u(C_t, H_t)$$

$$s.t.\ \omega F_t - \delta \leq \Omega_t$$

$$C_t + x_{Ft} + x_{Gt} + x_{Nt} = Y_t$$

Law of Motion of knowledge and initial knowledge stocks

$$L_{Ft} + L_{Gt} + L_{Nt} \leq H,$$

$$s_{Ft} + s_{Gt} \leq S.$$

Production of Y_t is defined by the equations describing production in the model. It holds that $x_{Jt} = \int_0^1 x_{Jit} di$. The variable PV stands in for the continuation value of the economy; see Online Appendix C for the derivation.

C Numerical appendix

Since I cannot solve explicitly for the optimal policy over an infinite horizon, I truncate the problem after period T. In the literature, utility in periods after T are approximated under the assumption that policy variables are fixed, and the economy reaches a balanced growth path (Barrage, 2020; Jones et al., 1993). However, assuming a constant carbon tax would most likely violate the emission limit since the model is designed to reflect market forces describing an economy with green and fossil sectors operating in equilibrium.

I motivate the design of the continuation value by assuming the planner would hand over the economy to a successor after period T. A continuation value, PV, in the objective function captures that the planner cares about utility after period T. This set-up accounts for concerns about economic well-being of future generations in a similar vein than the sustainability criterion proposed by the World Commission on Environment and Development (1987) by attaching some value to the final productivity level.[42] I

[42]The sustainable development criterion reads "*[...] to ensure that it meets the needs of the present without comprising the ability of future generations to meet their own needs.*" (p.24). This is a vague definition. Dasgupta (2021) p.(332) interprets this criterion as meaning: "*[...] each generation should bequeath to its successor at least as large a productive base as it had inherited from its predecessor.*" However, this cannot be used to derive a sensible condition on the optimization in the present setting since there is no negative growth and technology is the only asset bequeathed to future generations. Thus, successors will always have at least as much productive resources as predecessors left. The relation to the future is instead approximated by a future potential to derive utility from consumption given bequeathed technology levels. Natural needs of the future are accounted for through the emission limit.

approximate the value of future productivity by assuming constant growth rates. To mitigate concerns that the choice of the continuation value drives the results, I experiment with the exact value of explicit optimization periods. I truncate the problem once explicitly adding a further period leaves the optimal allocation numerically unchanged. That is the case after $T = 42$, or 210 years. The planner's objective function becomes:

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

In more detail, I define the continuation value as the consumption utility over the infinite horizon starting from the last explicit maximization period:

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

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. Then, consumption grows at the same rate as output. Second, as an approximation to future growth, I assume the economy grows at the same rate as in the last explicit optimization period. Let $\gamma_{yT} = \frac{Y_T}{Y_{T-1}} - 1$. Under above assumptions, I can rewrite future consumption as $C_s = (1 + \gamma_{yT})^{s-T} C_T$. Given the functional form

$$u(C_s) = \frac{C_s^{1-\theta}}{1-\theta},$$

the continuation value reduces to

$$PV = \beta^T \left(\frac{1}{1 - \beta(1+\gamma_{yT})^{1-\theta}} \frac{C_T^{1-\theta}}{1-\theta} \right).$$

D Calibration

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

D.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 within sectoral 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 US applicants[43] from the European Patent Office (EPO)'s database PATSTAT. I consider patents filed between 1950 to 2010.[44] To classify patents into the three

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

[44] The number of granted patents displays a sharp reduction after this year due to the time which may elapse from applying for protection to a patent being granted. The data on granted patents for more recent years is, thus, not complete.

distinct sectors of the model, I rely on the classification provided by a joint work of the International Energy Agency (IEA) and the EPO.[45].

Using the number of patents[46] related to green, fossil, and non-energy 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. More precisely, I calculate sector-specific knowledge, K_{Jt}, as:

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

where R_{Jt} stands in for new patents in sector J. 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 US 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.

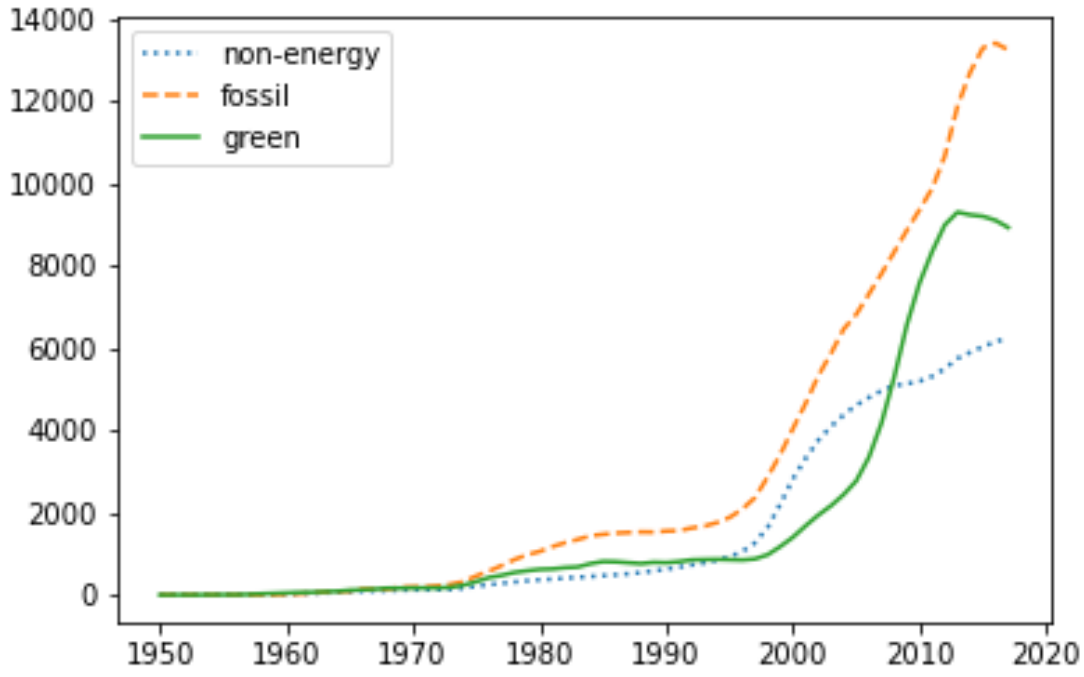
D.2 Measuring sectoral expertise

First, I derive scale-dependent measures of expertise based on historic cumulative output as described by the model's LOM for expertise, equation 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 me with a proxy for the ratio of green-to-fossil expertise. Jointly with the long-run ratio of knowledge stocks this provides the ratio of fossil-to-green productivities, A_f/A_g, on the BGP. Based on this information and the additional calibration steps that determines the non-research side of the model yield 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.

[45]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

[46]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 is flawed in that it depends on the structure of the economy and the green transition. 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, for example. 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.

Figure 5: Annual knowledge stock by sector



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 US 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.

D.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, total factor productivity in sector J in the model is given by, $A_J^{1-\alpha_J}$. Define $\iota_L\tilde{J}$ and $\iota_K\tilde{J}$ as $A_J^{1-\alpha_J} = Q_{Jt}^{\iota_L\tilde{J}} K_{Jt}^{\iota_K\tilde{J}}$.

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 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 total factor productivity 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: $\tilde{\iota_{LJ}} = \frac{\log(LR+1)}{\log(2)}$.

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[47], since semiconductors are used in a variety of sectors not necessarily related to energy. $\iota_{\tilde{LN}}$ is then backed out as $\iota_{\tilde{LN}} = -\frac{ln(1-0.2)}{ln(2)} = 0.32$.[48] Assuming that elasticities of labor-augmenting productivities are the same across sectors, we have that $\iota_L = \frac{\iota_{\tilde{LN}}}{1-\alpha_n} = 0.503$.[49]

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 percent increase in TFP. I match this measure to the non-energy sector yielding $\iota_K = 0.2148$ as the elasticity of labor-augmenting technology with respect to the knowledge stock.[50] 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.

[47]Even though growth in the non-energy sector is modeled 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.

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

[49]This results in the following elasticities of fossil and green total factor productivity 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.

[50]Sector-specific TFPs are given by $\iota_{\tilde{KG}} = 0.0302$,$\iota_{\tilde{KF}} = 0.0537$.