



IMF Working Paper

Research Department

Can International Technological Diffusion Substitute for Coordinated Global Policies to Mitigate Climate Change?**Prepared by Philip Barrett[†]**

Authorized for distribution by Malhar Nabar

June 2021

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

In short, yes. I use a multi-region integrated assessment model with fuel-specific endogenous technical change to examine the impact of Europe and China reducing emissions to zero by mid-century. Without international technological diffusion this is insufficient to avoid catastrophic climate change. But when innovation can diffuse overseas, long-run temperature increases are limited to 3 degrees. This occurs because policy not only encourages green innovations but also dissuades dirty innovations which would otherwise spread. The most effective policy package in emissions-reducing regions is a research subsidy funded by a carbon tax, driven in the short term by the direct effect of the carbon tax on the composition of energy, and later by innovation induced by research subsidies. Green production subsidies are ineffective because they undermine incentives for innovation.

JEL Classification Numbers: E60, F42, Q54

Keywords: Climate Change, Technology, Spillovers

Author's E-Mail Address: pbarrett@imf.org

[†] I thank Oya Celasun, Wenjie Chen, John Hassler, Florence Jaumotte, Per Krusell, and Philippe Wingender for their helpful comments and suggestions.

1 Introduction

The scientific consensus is clear: the risk of catastrophic climate change can only be avoided if global emissions of carbon dioxide fall almost to zero within the next few decades. This of course implies emissions must fall everywhere. This fact is often used to argue that coordinated global policy action is essential, on the grounds that emissions will only decline everywhere if mitigation policies are implemented everywhere. This paper investigates a counterpoint to this argument, asking whether international technical diffusion might mean that policy action in just a few major emitters could be enough to limit global emissions. That is: can international technological diffusion substitute for coordinated global policies to mitigate climate change?

This argument relies on two central ideas. First, that climate mitigation policies induce greener technological change. Emissions-reducing policies such as carbon taxes have an expenditure-switching effect, increasing the relative size of the market for low-emission energy technologies, such as solar panels and wind turbines. Larger markets then incentivize further innovation in these technologies, termed “induced technical change”. Second, that innovations in one place diffuse to others. In particular, policy-induced innovations in low-emissions technologies in one part of the world will spread to others. Together, these two mechanisms create a positive international policy spillover. By inducing green innovation which diffuses overseas, policies to reduce emissions in one place will lower emissions elsewhere, even in regions which do not pursue climate mitigation policies.

This positive international policy spillover seems particularly relevant at present. As of Spring 2021, major emitters – including the European Union, China, and the United States – have all announced ambitious plans to reduce emissions in the coming years, suggesting that a new era of international climate policy co-ordination is at hand. However, history is replete with sobering reminders that such international co-ordination is easy to promise but hard to deliver. For example, the Intergovernmental Panel on Climate Change (IPCC) was created in 1988. And the Kyoto protocol was intended to establish legally binding limits on individual countries’ greenhouse gases back in 1997. Yet emissions have continued to rise, increasing by more than 50 percent since 1990 and global temperatures have risen by almost 1 degree Celsius since 1980.¹

This suggests that we should take seriously the risk that not all major emitters will be able to meet their targets. Yet assessing the impact of this risk depends critically on the strength of international policy spillovers. If such spillovers are weak, then failure of one major emitter to implement policies which drive their emissions to zero will keep global emissions from hitting zero, even in the long run. But if spillovers are sufficiently strong, then policies elsewhere can act as a partial substitute for policies at home. In this case, new green technologies will diffuse from overseas, lowering the cost of clean energy and pulling down emissions even in parts of the world with weak climate change mitigation policies. Of course, even with strong spillovers, global policy co-ordination remains the first best remains. But the strength of international policy spillovers determine the cost of failing to attain the first best. This paper is an attempt to understand the magnitude of these policy spillovers and thus the extent to which they limit global losses when only some countries implement policies to

¹Given the slow dynamics of temperature this likely understates the seriousness of the current situation. Roughly another one degree of warming is likely already unavoidable, due to the delayed warming effect of past emissions.

reduce emissions.

To investigate this issue, I construct a regional integrated assessment model of the global economy and climate. Three key features make this framework well-suited to assessing the likely size of these international policy spillovers.

The first important feature is a standard climate externality. Producing energy from dirty fossil fuels emits greenhouse gases, which raise global temperatures and in turn reduces productivity. First popularized by the influential DICE model (see Nordhaus (1993) and more recently Nordhaus (2014a)), this externality is now standard in the literature on climate change mitigation.² This is a key source of inefficiency in the model and is the motive for climate change mitigation policies. As in Nordhaus (2010), the model is regional, with each region generating emissions and setting independent climate mitigation policies. And similar to Hassler et al. (2020), no region is sufficiently large as to prevent catastrophic climate change by simply reducing their own emissions (absent positive policy spillovers). The standard solution to this failing is coordinated and collective action to address climate change.

The second important feature is endogenous directed technical change. Energy producing firms can engage in fuel-specific research which lowers their costs of energy production using that fuel (or equivalently, increases their productivity). As first formulated in Acemoglu (2002), and more recently applied in climate-related settings by Acemoglu et al. (2016) and Fried (2018), incentives for research are a function of the total input cost, i.e. price times quantity, known as “market size”.³ Intuitively, the more an energy-producing firm spends on an input, the greater their demand for cost-reducing efficiency gains using that fuel. A byproduct of policies which reduce emissions is that they increase the relative market size of less-polluting energy inputs, shifting research towards these technologies and lowering the input cost of these fuels. This then drives the direction of technical change in the long run, in some cases permanently.⁴

The third important feature is international technological diffusion. At each point in time, the starting point for new research in a given region is a function of both past domestic and foreign technology. This means that innovations in one region slowly diffuse worldwide. As technologies are specific to fuel types (e.g. coal, solar, etc.), policies which affect research incentives in one region impact future relative costs of clean and dirty fuels not only domestically but also overseas as well. The international diffusion of technological advances is a relatively well-studied topic, so I draw on this literature when defining the functional form and speed of technological flows. Relevant papers include Eaton and Kortum (1999) – the genesis of most modern work on this topic – as well as Comin and Hobijn (2010) and Dechezleprêtre et al. (2013) which distinguish possible differences between diffusion of technologies in different sectors with differing carbon emissions respectively. More recent work, such as Caselli and Coleman (2001), Sampson (2016), Bloom et al. (2016), and Buera and Oberfield (2020), emphasizes the role of trade – particularly imports – in defining international technological

²More recent exponents of this mechanism include Golosov et al. (2014) and Barrage (2020).

³An earlier literature included various other endogenous mechanisms for technical change such as learning-by-doing, but but lacked a framework for comparing the relative economic effects of different policies on the direction of technical change. See Van der Zwaan et al. (2002), Buonanno et al. (2003), Bosetti et al. (2006), Popp et al. (2010), and a related critique in Nordhaus (2014c)

⁴See Acemoglu (2002), Acemoglu et al. (2012), and discussion in Appendix C of this paper.

flows, which I study in a robustness exercise.

Together, these features allow a quantitative assessment of the titular question of this paper. The particular policy experiment I study is one where OECD Europe and China commit to reducing emissions almost to zero by mid-century, in line with current policy plans.⁵ I consider three different ways that each region might meet their emissions targets: a carbon tax alone, green energy research subsidies funded by a carbon tax, and green energy production subsidies funded by a carbon tax. In each case, I invert the full model numerically to back out the policies which attain specific emissions paths in OECD Europe and China in general equilibrium. Three main results emerge from this experiment.

The most important result is that the size of potential positive technological spillovers is large. Without any global policy action, the model predicts runaway and catastrophic climate change, with increases in global temperatures in excess of 7 degrees Celsius above pre-industrial levels by the end of the next century.⁶ The direct effect (i.e. without technological diffusion) of reduced emissions in China and Europe is relatively small, reducing long-run temperatures by around 1.5 Celsius – slowing the effect of climate change but failing to prevent runaway increases in global temperatures. When the technological spillovers of policy are taken into account the impact is much larger, preventing runaway climate change and limiting increases to less than 3 degrees above pre-industrial levels. This still represents a significant warming of the globe and would undoubtedly impose sizable costs on society. However, the difference between the cases with and without international technological diffusion is qualitatively huge: the former is a future where the planet likely becomes uninhabitable towards the end of this century; the latter is one where climate change is painful but manageable. This result also has implications for the strategic calculus around global cooperation: individual incentives to participate are stronger when spillovers are larger, because individual actions have a larger positive effect on global temperatures.

The second main result is an explanation of the mechanism driving this result. The natural assumption is that international technological diffusion works by spurring innovation in new green technologies. Indeed, that was the intuition used in motivating this paper just a few paragraphs prior. However, a slightly different channel is just as important: creating disincentives for innovation in dirty technologies. This is clearest in the case of China. As China is by far the world’s largest consumer of coal – accounting for around half of global consumption – incentives in China to cut the cost of extracting energy from coal are enormous. By reducing the domestic market for coal, policies which shift Chinese energy demand away from coal reduce these incentives to innovate. As a result, these

⁵This is not to say that other parts of the world will not successfully implement emissions-reducing policies, or that these particular emitters will be able to meet their goals. Rather, given the difficulty achieving a global consensus for action in the last few decades, an imperfect global commitment to climate mitigation policies seems like a reasonable scenario to consider.

⁶A lack of precision over the exact values of important environmental parameters mean that long-run temperature predictions are hugely uncertain, so one should interpret the level of long-run temperatures as illustrative rather than precise. Instead, the important point is that the marginal effect is large. Because changes in key environmental parameters apply to both the baseline and the policy scenarios, the marginal effect remains similar as these vary. Nevertheless, the baseline case is in line with other laissez-faire scenarios in the literature, such as Golosov et al. (2014) and Hassler et al. (2020)

policies cause some cost-saving innovations in coal to *not* happen. By not existing, these innovations cannot diffuse to other regions. This is an essential part of the mechanism by which emissions-reducing policies have a positive policy spillover to other countries, dissuading socially harmful technological development, raising the cost of pollution and driving down emissions worldwide.

The third main result is to characterize what determines the policy mix needed to meet emissions targets in Europe and China. To do this I build a partial equilibrium approximation to the energy market in a closed economy. This allows an approximate but accurate decomposition of policy responses in the full model into three channels: two are static – the composition of energy inputs and total energy consumption, akin to substitution and income effects in a consumer problem – and one dynamic – induced technical change. These help explain why policies in emissions reducing regions work and allows a comparison of the efficacy of the different policy instruments. Amongst other things, I show that carbon taxes are essential in the short run because the two static channels reinforce and because the tax base (fossil fuels) is large. In contrast, subsidies to either green production or research are contemporaneously ineffective because the static channels offset and the subsidy base (green energy) is small. I show further that green research subsidies are preferable to production subsidies because they have a much larger dynamic effect. The reason is that production subsidies lower firms' input cost share of renewable energy, disincentivizing research and offsetting the induced technical change that would otherwise occur – why bother innovating if something is already cheap? As a result, a policy package which funds research subsidies is the most effective at reducing emissions and allows a reduction in medium run carbon taxes around three decades sooner (or, lower by around \$100 per tonne CO₂) than for a carbon tax alone; the equivalent effect for production subsidies is essentially zero. I also provide an explicit dynamic decomposition of the effect this policy package, showing that in the long run induced technical change is the dominant effect, split roughly 50-50 between the effects of the carbon tax and the research subsidy.

Of course, uncertainty over both economic and environmental parameters means that long run outcomes are exceptionally hard to predict. And so the specific outcomes resented in this paper should be thought of as scenarios rather than forecasts. That said, the model used here aims to capture a reasonable consensus view on the relationship between economic activity, international technological diffusion, and climate change – matching not only the key moments of the data but also important dynamic properties of other models. An important consequence of international policy spillovers not studied here in depth is their subsequent impact on incentives for participation in global climate agreements. This paper considers the impact of spillovers on a given global policy scenario. However, because spillovers affect the marginal gains of individual countries deciding to join or leave a climate agreement, the magnitude of these spillovers could affect participation in a global plan to tackle climate change. This impact could be positive (say, for a large region near the technological which can internalize the global benefit of the spillovers that their policy cooperation induces) or negative (for a region far from the technological frontier, which benefits from foreign technological gains whether they participate or not). Analyzing this with any degree of sophistication is beyond the scope of this paper. In any case, the first best remains global cooperation – something for which technological spillovers can only ever partially substitute.

The rest of the paper proceeds as follows. Section 2 outlines the model. Section 3 describes how the model is calibrated and presents the laissez-faire outcome. Section 4 discusses the domestic aspect of emissions reduction in Europe and China – the effectiveness of different policy packages in meeting the emissions target and the impact on economic outcomes. Section 5 then discusses the international spillovers of these policies and investigates the role of parametric assumptions about the speed of international diffusion of technology. Section 6 shows that these conclusions are results are robust to assumptions in the model about the form of endogenous technical change and determinants of the savings rate. Section 7 concludes.

2 Model

In this section I outline the model. It has multiple regions and features a feedback from economic activity to global temperature via emissions (known as an “integrated assessment model”). Emissions come from the production of energy, which is the overwhelming source of the gasses which cause climate change – most notably, carbon dioxide.⁷ Directed endogenous technical change plays a central role, with the greenness of technological developments responding to climate change mitigation policies. Because technological change in each region diffuses to the others over time, this produces an international policy externality.

The basics of the model follows closely Hassler et al. (2020). Each of M regions produces a domestically-consumed final output good which takes labor, capital, and energy as input. Energy is produced as a CES aggregate of N fuels in each region, with one fuel an internationally-traded exhaustible resource – conventionally-produced oil – and all others domestically produced. Emissions, which are a byproduct of consuming a subset of fuels, raise global temperatures, causing a climate externality. There are three main extensions relative to Hassler et al. (2020). The most important are the explicit inclusion of endogenous technical change⁸ and scope for international technological spillovers. The other extension is to include a non-unit elasticity of energy demand.

The remaining details of the model are kept as simple as possible, allowing a focus on the differentiated impacts of R&D in different regions, and their international spillovers. In particular, the model is set up to guarantee that the optimal savings rate of households is always constant. This is a deliberate result of the modeling choices and is needed because the subsequent policy experiments are rather complex: In Section 4 I back out the policies required to hit a dynamic emissions target in two regions. Doing this in a dynamic general equilibrium model with multiple regions (each with multiple policies), a common exhaustible resource, endogenous technical change, and a global climate externality is potentially quite challenging. By making modeling choices which guarantee that savings rate of households is constant, I can dramatically simplify the problem and invert the model period-by-period.⁹ In Section 6 I show that this simplification is not quantitatively important: feeding

⁷See Pachauri et al. (2014).

⁸Hassler et al. (2020) propose a similar form of endogenous technical change, with unit returns to scale in R&D. This means that interior solutions only exist for complementary fuel types, which is a likely counterfactual parameter restriction. They also do not solve the general equilibrium solution, instead using an ad hoc modification to exogenous trends in fuel-specific technologies to proxy for this channel.

⁹In contrast, a time-varying savings rate would necessitate simultaneous search across policies, regions, *and* time

the policies which hit emissions targets in the constant-savings-rate model into a richer model with time-varying savings generates very similar aggregate outcomes.

2.1 Aggregate production

Time is discrete and indexed by t . In each of $m = 1, \dots, M - 1$ regions, a continuum of competitive firms produce a final good using a production technology which aggregates energy and a Cobb-Douglas capital-labor bundle with a constant elasticity of substitution.¹⁰ Gross output is then given by:

$$Y_{m,t} = z_{m,t} \left((1 - \nu_m) Y_{m,t}^*{}^\sigma + \nu_m E_{m,t}^\sigma \right)^{\frac{1}{\sigma}} \quad (1)$$

Where $Y_{m,t}^* = (A_{m,t} L_{m,t})^{1-\alpha_m} K_{m,t}^{\alpha_m}$ is the Cobb-Douglas labor-capital input mixture, $E_{m,t}$ is the intermediate energy input, $L_{m,t}$ is labor, $K_{m,t}$ is capital, $A_{m,t}$ exogenous labor-augmenting aggregate technical change, and $z_{m,t}$ is a region-specific aggregate productivity loss due to climate change (discussed later). The energy and labor share parameters ν_m and α_m are region-specific and will be identified from the data. I use $F = 1/(1 - \sigma)$ to denote the elasticity of substitution between the capital-labor bundle and energy.

Final goods firms rent capital and labor from households, and purchase energy from energy-producing firms. The respective factor prices $i_{m,t}$, $w_{m,t}$, and $P_{m,t}$ are determined by the first order conditions:

$$i_{m,t} = \alpha_m (1 - \nu_m) z_{m,t}^\sigma \left(\frac{Y_{m,t}^*}{K_{m,t}} \right) \left(\frac{Y_{m,t}}{Y_{m,t}^*} \right)^{1-\sigma} \quad (2)$$

$$w_{m,t} = (1 - \alpha_m) (1 - \nu_m) z_{m,t}^\sigma \left(\frac{Y_{m,t}^*}{L_{m,t}} \right) \left(\frac{Y_{m,t}}{Y_{m,t}^*} \right)^{1-\sigma} \quad (3)$$

$$P_{m,t} = \nu_m z_{m,t}^\sigma \left(\frac{Y_{m,t}}{E_{m,t}} \right)^{1-\sigma} \quad (4)$$

2.2 Energy production and research technologies

Each period, energy-producing firms have access to technologies for production of energy and research. They use the energy technology to combine fuels into energy, and can conduct research to lower the cost of producing energy from a given fuel type.

Fuels: There are N fuels, which are the inputs into energy production. Fuels differ by their emissions intensity and production costs. Fuel 1 is conventional oil, of which there is a fixed stock, owned by region M . Fuels $i = 2, \dots, N$ are produced with a linear technology using $l_{i,t}$ units of the final output good to produce $e_{m,i,t}$ units of energy,

$$e_{m,i,t} = x_{m,i,t} l_{m,i,t}$$

periods. In my headline experiment, this is a root-finding problem in 136 dimensions. By solving period-by-period I can convert this to a sequence of four-dimensional problems.

¹⁰Region M is special. It produces oil. See discussion below.

where $x_{m,i,t}$ is the fuel-specific technology.

Emissions intensity is constant for each fuel and denoted g_i , so emissions for fuel i are given by:

$$m_{m,i,t} = g_i e_{mi,t}$$

In terms of emissions, one fuel is special. Fuel j is a zero-carbon renewable energy source, i.e. “green” energy. Thus, $g_j = 0$.¹¹ Fuels are best interpreted as a particular form of energy with a specific emissions intensity. For example, gallons of oil, or kilowatt-hours of solar electricity.¹² As $l_{i,t}$ is expressed in units of the numeraire, the (average and marginal) cost of producing fuel i is $p_{i,t} = 1/x_{i,t}$.

An important assumption here is that for a given level of technology, each fuel except oil is available at constant input cost. This is identical to Hassler et al. (2020) and equivalent to assuming that each fuel is in infinite supply. For renewable energy, this is clearly a reasonable assumption. The supply of sunlight and wind are both effectively unlimited, and sites for generators remain abundant (hydroelectric power is perhaps an exception). Only the extent of technological capabilities limits the marginal cost of collecting this energy source. In contrast, dirty fuels are in reality in fixed supply, so the possibility of their exhaustion, and associated increase in their prices, could in theory act to mitigate climate change within the foreseeable future. However, the stocks of most dirty fuels are so large that their scarcity value only becomes relevant in scenarios far outside the range of possibilities that we will discuss. For instance, coal reserves will only run out once the global temperature has risen well beyond any tolerable levels. The only major dirty fuel type for which this is not true is conventionally-produced oil. Stocks of conventionally-produced oil are sufficiently small that their exhaustion is possible within the next century. Thus, I explicitly model the exhaustion of this stock but treat others as being in infinite supply – a simplification for sure, but not an unreasonable one.¹³

Research: For fuels $i = 2, \dots, N$, energy-producing firms can conduct research to lower the cost of production.¹⁴ As in Fried (2018) I assume that the period length is sufficiently long that firms can only capture the gains from research within the current period, either because formal property rights (e.g. patents) expire or because competitors can backwards-engineer close substitutes.

The reference level of technology $\bar{x}_{m,i,t}$ captures the extent of pre-existing knowledge, and the cost of new learning, $r_{m,i,t}$, is expressed relative to this level:

¹¹This does not mean that fuel j is the *unique* zero-carbon fuel. Other energy source, most notably nuclear fuels, are also zero carbon but we typically think of them as distinct from the class of “green” fuels.

¹²I return to the issue of units in Section 3.1.

¹³Some illustrative calculations: global conventional oil stocks are some 204 Gigatonnes but Coal around 1054 Gigatonnes (see OPEC Annual Statistical Bulletin 2019 and BP Statistical Review of World Energy 2019), implying that coal and oil sequester around 2740 and 510 Gigatonnes of carbon each. Given that around 20 percent of emissions remain in the atmosphere permanently, using these sources in their entirety would imply long-run increases in the stock of global carbon in the order of 550 and 102 Gigatonnes. Given that the total increase in atmospheric carbon since the start of the industrial era is around 200 Gigatonnes, any scenario in which global coal stocks are exhausted would leave the climate so changed that mitigation policies would be an irrelevance.

¹⁴Mirroring Hassler et al. (2020), fuel $i = 1$ is conventional oil, an exhaustible resource produced by region M . So that the input price of oil maps directly to global oil prices, I assume that the oil production technology is unimprovable.

$$\begin{aligned}
r_{m,i,t} &= \frac{\epsilon_{m,i}(1 - \chi_{m,i,t})}{\eta - 1} \left(\frac{x_{m,i,t}}{\bar{x}_{m,i,t}} \right)^{\eta-1} \\
&= \frac{\epsilon_{m,i}(1 - \chi_{m,i,t})}{\eta - 1} \left(\frac{p_{m,i,t}}{\bar{p}_{m,i,t}} \right)^{1-\eta}
\end{aligned} \tag{5}$$

where $\eta > 1$, $\epsilon_{m,i}$ is a fuel-specific cost parameter, and $\chi_{m,i,t}$ is a government subsidy to research in fuel i .¹⁵ For simplicity I assume that research subsidies do not vary, except for the green technology. Without loss of generality, I therefore set $\chi_{i,t} = 0$ for $i \neq j$.¹⁶ The return to scale in energy research is thus $1/(\eta - 1)$. It is not costless to hold technology at the reference level $\bar{x}_{i,t}$ which one can think of this as the cost of maintaining sufficient human capital to keep the current level of technology fixed. This specification is a generalization of the log research function considered by Hassler et al. (2020), which we recover in the limit as $\eta \rightarrow 1$.

Energy: Energy-producing firms produce fuels and bundle them into energy. The production function is CES with elasticity $\mathcal{E} = 1/(1 - \rho)$.

$$E_{m,t} = \left(\sum_{i=1}^n \lambda_{m,i} e_{m,i,t}^\rho \right)^{1/\rho} \tag{6}$$

2.3 Energy supply

Climate mitigation policies (carbon taxes, production subsidies, and R&D subsidies) are region-specific and collected from the energy producing firm. The input cost function for energy-producing firms is given by:

$$\begin{aligned}
c(e_{1,m,t}, \dots, e_{N,m,t}, p_{1,t}, p_{2,m,t}, \dots, p_{N,m,t}) &= (\tau_{m,t} g_{1,m} + p_{1,m,t}) e_{i,m,t} \\
&+ \sum_{i=2}^n \left((\tau_{m,t} g_{i,m} + (1 - \delta_{i,m,t}) p_{i,m,t}) e_{i,m,t} + \frac{\epsilon_{i,m}(1 - \chi_{i,m,t})}{\eta - 1} \left(\frac{p_{i,m,t}}{\bar{p}_{i,m,t}} \right)^{1-\eta} \right)
\end{aligned} \tag{7}$$

Where $\bar{p}_{i,w,t}$ is the technologically-determined reference cost of production of fuel i in region m , and $\delta_{i,m,t}$ and $\chi_{i,m,t}$ are the purchasing and R&D subsidies for fuel i in region m , which I assume are zero except for $i = j$.

At any level of energy production, profit-maximizing firms will always minimize their costs. Thus,

¹⁵The idea here is that the government picks up a share $\chi_{m,i,t}$ of the bill for research, meaning that private research expenditure, $r_{m,i,t}$ is a fraction $1 - \chi_{m,i,t}$ of total funding for research. Government subsidy expenditure is thus $(\chi_{m,i,t}/(1 - \chi_{m,i,t})) \times r_{m,i,t}$

¹⁶Note that a constant subsidy for any fuel can be captured by the level of ϵ_i .

the output cost function for the energy-producing firm is:

$$C(E) = \min_{\{e_{1,m,t}, \dots, e_{N,m,t}, p_{1,t}, p_{2,m,t}, \dots, p_{N,m,t}\}} c(e_{1,m,t}, \dots, e_{N,m,t}, p_{1,t}, p_{2,m,t}, \dots, p_{N,m,t})$$

$$\text{s.t. } \left(\sum_{i=1}^n \lambda_{m,i} e_{m,i,t}^\rho \right)^{1/\rho} = E$$

If $\mathcal{E} < \eta$, this problem has an interior solution with first-order conditions:

$$\frac{e_{i,m,t}}{e_{j,m,t}} = \left(\frac{\lambda_j \hat{p}_{i,m,t}}{\lambda_i \hat{p}_{j,m,t}} \right)^{-\mathcal{E}} \quad \forall i \quad (8)$$

$$(1 - \delta_{i,m,t}) e_{i,m,t} = \frac{\epsilon_i (1 - \chi_{i,m,t})}{\bar{p}_{i,m,t}} \left(\frac{p_{i,m,t}}{\bar{p}_{m,i,t}} \right)^{-\eta} \quad \forall i > 1 \quad (9)$$

Where $\hat{p}_{i,m,t} = \tau_{m,t} g_i + (1 - \delta_{i,m,t}) p_{i,m,t}$ is the cost of production of fuel i inclusive of tax and subsidies.

If $\mathcal{E} \geq \eta$, then the cost-minimization problem does not have an interior solution. The intuition is that if the returns to scale in research or the intra-fuel elasticity of substitution are too high the research pays for itself. A firm can conduct research in just one fuel type, lowering the overall cost of production if it is sufficiently easy to substitute into the cheaper fuel.

Rearranging equation 9 it is straightforward to show that $r_{i,m,t} \propto e_{i,m,t} p_{i,m,t}$. That is, research spending is proportionate to expenditure on a given input. This is intuitive: because research lowers input costs, it makes sense to spend more on inputs which account for a larger share of firms' expenditures. This is analogous to the "market size" effect which was first proposed by Acemoglu (2002) and is common in the literature. This captures a similar idea, that because innovation is a fixed cost, potential profits for innovators (and hence incentives for innovation) are higher in markets where potential revenues are larger.

The resulting output cost function $C(E)$ is downward-sloping and so the energy market is a natural monopoly. I show this formally in Appendix A, but a replication argument can convey the point with less algebra. As research is a fixed cost, a firm can scale up production at the same input prices without changing research expenditure. But because the research cost is defrayed over more units, the average cost falls. I also show that the monopoly supply of energy is interior only if 1) the elasticity of demand is less than the inverse returns to scale in research (i.e. $F = 1/(1 - \sigma) < \eta$) and 2) that the monopolist's price is capped by regulation. I assume that this price cap sets monopoly profits to zero. This is the constrained first-best outcome under the constraint that the government does not provide a subsidy to energy production or consumption. In this case, price equals average cost. That is:

$$P_{m,t} = \frac{C(E_{m,t})}{E_{m,t}} \quad (10)$$

Note that the relative input shares are determined by the first-order conditions (8) and (9), which are identical to those which would arise if the energy market were competitive. In other words, it is only the quantity of energy supplied which is dependent on assumptions about how monopoly energy supply is regulated; the composition of energy is independent of these.

More broadly, the set-up of energy production features a number of simplifications relative to the literature, most notably the functional form of the research cost function but also in other aspects of energy production. These help keep the model lightweight enough that one can back out policies which meet specific emissions targets. However, the reader may quite naturally have questions or concerns about the importance of these modeling simplifications. These are addressed in Section 6, where I show compare the policy response of this model to two other cutting-edge frameworks, showing that these simplifications do not quantitatively affect the responses of emissions and energy usage to policies.

2.4 Governments

Governments in regions $m = 1, \dots, M-1$ raise revenues from the carbon tax, and spend on subsidies for production and R&D. For simplicity, I assume they run balanced budgets, rebating any excess revenue from the carbon tax over expenditure on subsidies as a lump sum transfer $\Omega_{m,t}$ (when $\Omega_{m,t} < 0$, a lump-sum tax).

The government's period budget constraint is:

$$\Omega_{m,t} + \delta_{j,m,y} p_{j,m,t} e_{j,m,t} + \frac{\chi_{j,m,t}}{1 - \chi_{j,m,t}} r_{j,m,t} = \sum_{i=1}^N \tau_{m,t} g_{i,m} e_{i,m,t} \quad (11)$$

2.5 Households

Households in regions $m = 1, \dots, M-1$ sell labor and rent capital to firms. They receive lump transfers $\Omega_{m,t}$ from the government. Labor supply is inelastic and of size $L_{m,t}$. Capital depreciates fully each period and preferences are logarithmic over consumption. The household's problem is thus:

$$\begin{aligned} & \max \sum_{t=0}^{\infty} \beta^t \log C_{m,t} \\ \text{s.t.} \quad & C_{m,t} + K_{m,t+1} = w_{m,t} L_{m,t} + i_{m,t} K_{m,t} + \Omega_{m,t} \end{aligned}$$

The Euler equation is then:

$$\frac{C_{m,t+1}}{C_{m,t}} = \beta i_{m,t+1} \quad (12)$$

Defining factor income $\hat{Y}_{m,t} = w_{m,t} L_{m,t} + i_{m,t} K_{m,t}$, the ratio of transfers to factor income by $\Gamma_{m,t} = \Omega_{m,t} / \hat{Y}_{m,t}$, and the savings rate as $s_{m,t} = K_{m,t+1} / ((1 + \Gamma_{m,t}) \hat{Y}_{m,t})$, this simplifies to.

$$\frac{1 - s_{m,t+1}}{1 - s_{m,t}} = \frac{\alpha \beta}{s_{m,t} (1 + \Gamma_{m,t+1})} \quad (13)$$

Note that if $\Gamma_{m,t} = 0$ then the savings rate is constant at $s_{m,t} = \alpha \beta$. As is well known, a standard growth model with log preferences, Cobb-Douglas production and full depreciation of capital will deliver exactly this savings rate. But production is not Cobb-Douglas; rather, it is CES. So how does this happen and why is Γ_t relevant? The important point is that the standard result arises

not because aggregate production is Cobb-Douglas per se, but that household incomes come from payments to factors employed in a Cobb-Douglas production function. Because the input to energy production and research is the final good, rather than remunerated factor of production, households receive no factor income from payments to energy. Their income comes only from payments to labor and capital, which are factors in a Cobb-Douglas production function, as required. The exception is when $\Omega_t \neq 0$. Because government taxes and subsidies are covered by payments to energy-producing firms, government transfers are effectively partial factor payments to households from energy. In other words, households derive some income from the non-Cobb-Douglas part of the production technology. This also means that apparently simple static extensions, such as using household labor for energy production or research, would have an important dynamic implication – a non-constant savings rate.

2.6 Trade in conventional oil

Production of conventional oil follows Hassler et al. (2020) identically. Region M is assumed to contain the global stock of conventional oil, R_t . This is exchanged for final goods from other regions at the global oil price $p_{1,t}$, which are then consumed. Hassler et al. (2020) show that if households in the oil-producing region have logarithmic preferences over consumption, then the oil-producing region sells a fraction $(1 - \beta)$ of the stock of oil each period, consuming the proceeds.

$$R_{t+1} = \beta R_t \tag{14}$$

$$C_{M,t} = p_{1,t}(1 - \beta)R_t \tag{15}$$

The global market clearing condition for oil is then

$$\sum_{m=1}^{M-1} e_{1,t} = (1 - \beta)R_t \tag{16}$$

2.7 Technological diffusion

Technological diffusion occurs within a country in two ways. First, within-fuel inter-temporal knowledge spillovers. Progress in a given fuel technology enters into the public knowledge with a lag, and is freely available to future energy-producing firms. The second is intra-fuel spillovers. This is the idea that developments in one technology could have applications in others, such as when low-carbon technologies reduce the costs of dirtier technologies.¹⁷ Intra-fuel spillovers matter because they prevent one fuel from attaining an insurmountable technological lead. If intra-fuel spillovers are small, the input share of the most efficient fuel will increase in the next period. This stimulates further research and drives down costs of production, cementing the technological advantage of that fuel. This is exactly the intuition behind the mechanism in Acemoglu et al. (2016), who show that temporary climate mitigation policies can lead to permanent changes in fuel shares in the long run. With technological spillovers, advances in one fuel also pull down costs of production in others, preventing runaway progress by one type of fuel.

¹⁷Fried (2018) cites the example of innovation in solar cells being used to provide light for oil rigs.

Let $p_{m,i,t+1}^d$ be the input cost of fuel i in region m excluding imported technologies. I assume that this follows

$$\frac{p_{m,i,t+1}^d}{\bar{p}_{m,i,t}} = \left(\prod_k \left(\frac{p_{m,k,t}}{\bar{p}_{m,k,t}} \right)^{\rho_{ik}} \right)^\omega \quad (17)$$

That is, the increase in public knowledge about fuels i is a geometric weighted average of the rates of increase in each of the other technologies, with weights given by the $\rho_{i,k}$. The ρ_{ij} thus determine the degree of inter-fuel technological spillovers. We can express these more succinctly as the matrix ρ , the rows of which sum to unity. The parameter ω is a dampening parameter, affecting the aggregate accumulation of technology. If some part of innovation is technology cannot be passed on – perhaps because innovations are particularly well-adapted to a specific time and place – then one would set $\omega < 1$.

We can rewrite equation (17) more conveniently in logs.

$$\log \bar{\mathbf{p}}_{m,t+1}^d = (I_N - \omega\rho) \log \bar{\mathbf{p}}_{m,t} + \omega\rho \log \mathbf{p}_{m,t} \quad (18)$$

Where $\mathbf{p}_{m,t} = (p_{1,m,t}, \dots, p_{N,m,t})'$ and so on.

Beyond the two types of domestic technological diffusion, there is also a third – international technological diffusion, which is the main focus of this paper. This is modeled as incomplete catch-up to international technology frontier. The vector of global fuel-specific frontier technologies is $\bar{\mathbf{p}}_t^* = (\bar{p}_{2,t}^*, \dots, \bar{p}_{2,t}^*)$, where:

$$\bar{p}_{i,t}^* = \min_{m=1, \dots, M-1} \bar{p}_{m,i,t}^d \quad \forall i \quad (19)$$

Then the reference cost of production in each region is then a weighted average of domestic technological progress and catch-up to the international frontier.

$$\log \bar{\mathbf{p}}_{m,t+1} = \zeta_m \log \bar{\mathbf{p}}_t^* + (1 - \zeta_m) \log \mathbf{p}_{m,t}^d \quad (20)$$

where ζ_m measures the region-specific speed of diffusion of advances in the global technological frontier. If $\zeta_m = 1$, learning from foreign advances is immediate, whereas if $\zeta_m = 0$ it is so slow as to never occur. For simplicity I outline only the case where ζ_m is a single number, but one can easily extend this to allow for fuel-specific speeds of technological diffusion.

To summarize, the dynamics of technological change have both a domestic and an international component. The domestic component is governed by ρ , the matrix of intra-fuel spillovers, and ω , which determines the aggregate speed at which new technology becomes public. The international component is that a fixed fraction of foreign knowledge filters through to each region, with the speed varying region and measured by ζ_m .

2.8 The climate externality

The stock of atmospheric carbon S_t is a function of lagged global emissions $M_t = \sum_{m=1}^M M_{m,t}$. Following Golosov et al. (2014), a fraction φ_L of emissions are assumed to remain in the atmosphere permanently. Of what remains, a fraction $1 - \varphi_0$ leaves the atmosphere immediately and the remain-

der decays at rate φ .

$$S_t = \sum_{u=0}^{\infty} (\varphi_L + (1 - \varphi_L)\varphi_0(1 - \varphi)^u) M_{t-u} \quad (21)$$

This is also usefully written as:

$$\begin{aligned} S_t &= S_t^{perm} + S_t^{temp} \\ \text{where: } S_t^{perm} &= S_{t-1}^{perm} + \phi_L M_t \\ S_t^{temp} &= (1 - \phi)S_{t-1}^{temp} + (1 - \phi_L)\phi_0 M_t \end{aligned}$$

The stock of atmospheric carbon drives global temperatures. The dynamics of global atmospheric and oceanic temperatures are modeled as a coupled system. As in Hassler et al. (2020), this follows the linear energy budget model in DICE/RICE. Denoting by T_t global average surface temperatures, and by T_t^{ocean} average deep oceanic temperatures, the relationship between the carbon stock and temperature dynamics is given by:

$$\begin{bmatrix} T_t \\ T_t^{ocean} \end{bmatrix} = \begin{bmatrix} T_{t-1} \\ T_{t-1}^{ocean} \end{bmatrix} + \begin{bmatrix} -\sigma_1 \left(\frac{\eta^c}{\lambda^c} + \sigma_2 \right) & \sigma_1 \sigma_2 \\ \sigma_3 & \sigma_3 \end{bmatrix} \begin{bmatrix} T_{t-1} \\ T_{t-1}^{ocean} \end{bmatrix} + \begin{bmatrix} \sigma_1 \frac{\eta^c}{\log 2} \log \left(\frac{S_{t-1}}{\bar{S}} \right) \\ 0 \end{bmatrix} \quad (22)$$

Where \bar{S} is the pre-industrial stock of atmospheric carbon and λ^c is the climate sensitivity. A doubling of the stock of atmospheric carbon relative to pre-industrial levels will raise global temperatures by λ^c .

The temperature drives a region-specific climate externality which affects aggregate production via equation (1). As is standard, productivity declines are region-specific and quadratic in atmospheric temperature:

$$\log z_{m,t} = -(\phi_0^m + \phi_1^m T_{t-1} + \phi_2^m (T_{t-1})^2) \quad (23)$$

2.9 Equilibrium

The model is closed with the aggregate resource constraint:

$$C_{m,t} + K_{m,t} + P_{m,t}E_{m,t} = Y_{m,t} \quad (24)$$

Given region-specific policies $\{\tau_{m,t}, \delta_{j,m,t}, \chi_{j,m,t}\}$ and initial values $\{K_{m,0}, \bar{\mathbf{p}}_{m,0}\}$, global initial values $\{S_0, T_0, R_0\}$, and exogenous variables $\{A_{m,t}, L_{m,t}\}$, equilibrium is a sequence of global $\{S_t, T_t, R_t, p_{1,t}, \{p_{i,t}^*\}_{i=2}^N\}$ and region-specific variables $\{Y_{m,t}, E_{m,t}, C_{m,t}, K_{m,t+1}, \Omega_{m,t}, s_{m,t}, \Gamma_{m,t+1}, i_{m,t}, w_{m,t}, P_{m,t}, \{e_{i,m,t}, p_{i,m,t}, p_{i,m,t}^d, \bar{p}_{i,m,t}, r_{i,m,t}\}_{i=2}^N\}$ satisfying equations (1)-(6), (8)-(12), and (14)-(24).

3 Calibration and baseline solution

I calibrate the model to have $M = 12$ regions. Region M is the oil-producing region, for which I use OPEC countries. The other model regions are Asia-Pacific, China, European OECD, India, Japan, Latin America and the Caribbean, other OECD, Russia, Sub-Saharan Africa, the USA, and Rest of the World. I allow for $N = 7$ fuel types: conventional oil (imported from OPEC), coal, natural gas, unconventional oil (e.g. from fracking, produced domestically), hydroelectric power, nuclear power, and renewable fuels (mostly solar and wind, but also some biomass). As in Fried (2018), the time period in the model is five years. Period 0 corresponds to 2015-2019 so when I fit to period 0 values, the data uses the average of this time period for flow variables, and the level in 2015 for the stock.

3.1 Model calibration

Data sources. The main data source for the energy section is IEA World Energy Balances. This provides detailed country-year data on fuel-specific energy inputs and their associated carbon dioxide emissions, aggregate energy consumption (which differs from the supply of energy due to losses in production and transmission of energy), and some data on input prices. Fuel consumption is expressed in tonnes oil equivalent, a unit of energy, in line with the interpretation of $e_{i,m,t}$. This also guarantees that conversion of fuels with different energy intensities (e.g. types of coal or oil) are aggregated appropriately. Total emissions from fossil fuel combustion averaging a little around 31 Gigatonnes CO2 per year in the initial period, all from fossil fuel combustion. This captures the overwhelming majority of CO2 emissions.¹⁸ This data set includes some information on input prices, although coverage is spotty, mostly focused on advanced economies. For regions with no price data, I use the global average price for each fuel. The IEA data informs many of the region-specific parameters for the energy sector.

Exogenous variables. Aggregate labor productivity and labor forces in each region are exogenous. I use ILO labor force data to set the period 0 labor force, and ILO labor force projections to set the growth rate in the first period. Thereafter, the growth rate of the labor force in all regions converges smoothly to a long-run growth rate of zero (consistent with long-run UN population projections) at a rate which guarantees that global population growth in 2070 is 0.2% (also a UN projection). In line with Hassler et al. (2020), I assume that labor productivity in OECD countries grows at 0.75% per year, and that non-OECD countries converge to a growth path that grows at the same rate but is 60 percent above current levels. Convergence occurs at a rate of 25 percent per decade. That is:

$$\log A_{m,t} = \begin{cases} \log A_{m,t-1} + \log(1 + 5 \times 0.0075) & m \text{ is in OECD} \\ \log A_{m,t-1} + \log(1 + 5 \times 0.0075) + 0.25(\log A_{m,t-1} - \log \bar{A}_{m,t-1}) & m \text{ not in OECD} \end{cases}$$

Where: $\log \bar{A}_{m,t} = \log 1.6 + \log A_{m,0} + \log(1 + 5 \times 0.0075)t$

¹⁸Other important sources of atmospheric carbon dioxide are cement production and land use. The former accounts for around another 2 Gigatonnes of CO2 annually, and the latter in the order of 4 Gigatonnes although this is much more uncertain. See Olivier et al. (2017).

Common global parameters. Although some economic parameters vary by country, others are held fixed. The elasticity of intra-fuel substitution is taken from the mid-point of the estimates of Papageorgiou et al. (2017). Van der Werf (2008) provides estimates of the elasticity of substitution of energy with a Cobb-Douglas capital-labor bundle across countries and industries, from which I take the mean value. The concavity of the research production function η matches the returns to scale in Acemoglu et al. (2016), which here is given by $1/(\eta - 1) \simeq 3.6$ percent. This other key parameter choices are summarized in Table 1.

Cross-fuel spillovers are assumed to be symmetric, with a fraction $\bar{\rho}$ of future cost savings in any one fuel type attributable to past research in the same fuel, and a fraction $1 - \bar{\rho}$ coming from the other $N - 1$ fuels collectively. The spillover matrix ρ thus has elements:

$$\rho_{ij} = \begin{cases} \bar{\rho} & i = j \\ 1 - \bar{\rho}/(N - 1) & i \neq j \end{cases}$$

The value of $\bar{\rho}$ is set to 0.94, consistent with the estimates of Fried (2018).

For international technological spillovers ζ_m , I pick a baseline value of 0.45 for all regions and technologies, in line with the seminal study on this topic, Eaton and Kortum (1999). They estimate that the average lag between innovation and international adoption is around 11 years, implying a five-year diffusion rate of 0.45. This is also consistent with Comin and Hobijn (2010), who estimate that adoption lags for multiple new technologies, finding that those developed in the second half of the 20th century are adopted between 5 and 16 years later than in countries where they are first invented. A specific concern here might be that adoption of climate mitigation technologies are somehow different and that using general measures of adoption speed is as a result inappropriate. Fortunately, Dechezleprêtre et al. (2013) address exactly this question, finding no difference between climate mitigation and other technologies. However, because ζ_m is such an important parameter for this paper’s titular question, I later explore the sensitivity of my main results to changes in this parameter.

To match the long-run decline in global energy intensity of output, I assume an exogenous secular decline in the energy share parameter ν , replacing it with $\nu_{m,t} = (1 + g_\nu)\nu_{m,t-1}$. Given the steady annual decline in global energy intensity of around 1.5 percent per year since 1970, I assume that ν decreases at a similar rate.

Climate physics. The parameters for the dynamic climate model are all taken directly from Hassler et al. (2020), adjusting for differences in time period (five years here versus ten there). The long-run climate sensitivity, λ^c is a crucial determinant of the long-run impact of climate change on global temperatures. The meta-analysis of Meinshausen et al. (2009) estimates a median across a wide range of climate models of 3, which I match here.

The initial and pre-industrial stocks of atmospheric carbon are set to standard values in the literature, as are initial temperatures. The initial stock of conventional oil at the start of period 0 is taken from proven reserves of conventional oil, which are published annually by OPEC. This indirectly pins down β , via equation (16); international conventional oil demand must equal the fraction of the

oil stock consumed by the oil-producing region. The resultant β is equivalent to a subjective annual discount rate of 0.9 percent.

Object	Value	Rationale
<i>Economic parameters</i>		
Intra-fuel elasticity of substitution, \mathcal{E}	2.00	Mid-point of Papageorgiou et al. (2017)
Elasticity of energy demand, F	0.35	Consistent with Van der Werf (2008)
Research function parameter, η	28.88	Matches Acemoglu et al. (2016) at avg. gth.
Cross-fuel spillovers, $\bar{\rho}$	0.94	Matches Fried (2018) spillovers at avg. gth.
Persistence of innovation, ω	1.00	Fried (2018)
Rate of international technological diffusion, ζ	0.45	Eaton and Kortum (1999)
Discount rate, β	0.96	Matches global conventional oil supply
Change in energy share (percent annual), g_ν	-0.02	Matches global energy share since 1970
<i>Climate parameters</i>		
Persistence of atmospheric temp., σ_1	0.22	Hassler et al. (2020)
Convergence of air to ocean temp., σ_2	0.15	Hassler et al. (2020)
Permanent emissions share, σ_3	0.02	Hassler et al. (2020)
Temporary emissions share, ϕ_L	0.20	Golosov et al. (2014)
Emissions decay rate, ϕ_0	0.40	Golosov et al. (2014)
Convergence of ocean to air temp., ϕ	0.01	Golosov et al. (2014)
Long run climate sensitivity, λ^c	3.00	Meinshausen et al. (2009)
Short run climate sensitivity, η^c	1.85	Hassler et al. (2020)
Pre-industrial carbon stock (Gigatonnes), \bar{S}	581.00	Hassler et al. (2020)
<i>Initial values</i>		
Carbon stock (Gigatonnes), S_0	800.00	Golosov et al. (2014)
Permanent share of carbon stock, S_0^{perm}/S_0	0.18	Hassler et al. (2020)
Conventional oil stock (Gigatonnes), R_0	203.67	OPEC annual statistical bulletin, 2015
Atmospheric temp. vs. pre-industrial avg., T_0	0.85	Hassler et al. (2020)
Oceanic temp. vs. pre-industrial avg., T_0^{ocean}	0.01	Hassler et al. (2020)

Table 1: Calibration of common and global parameters and initial conditions

Region-specific parameters The remaining parameters are region-specific, matching regional aggregates and fuel-specific moments. As is standard, the energy share (taken from the IEA) and capital share (from the IMF’s World Economic Outlook database) pin down ν_m and α_m respectively. Where data are unavailable, as is the case for the capital share, I use the global average. The costs of climate change, which define the impact of the climate externality, are taken from the closest analogous regions in Nordhaus (2010).

As production and research in the energy sector are jointly determined, so too are the production and research parameters. Given an initial fuel-specific reference cost of production $\bar{p}_{m,i,0}$, the production weights $\lambda_{m,i}$ and research cost shifters $\epsilon_{m,i}$ are chosen to match input prices and fuel usages in the IEA data. Price data are expressed relative to a unit of energy – dollars per ton oil equivalent. That is, they measure the cost of using a given fuel type to produce the same energy as could be captured from burning one ton of oil. Prices are thus true input costs, consistent with the model.

Given this method for choosing production parameters to match input costs and quantities, the initial reference cost of production $\bar{p}_{m,i,0}$ thus determines the rate of technological improvement within

each fuel. Fuel-specific data on efficiency seem to be available only for the United States, from the Energy Information Administration’s Monthly Energy Review (EIA-MER). I thus pick $\bar{p}_{m,i,0}$ so that period-0 fuel-specific input cost reductions in the model equal the efficiency gains from EIA-MER over 2015-2019. The implicit assumption here is that productivity gains in different regions have progressed at approximately similar rates, even if the levels of the input costs are different. Of course, one could improve on this assumption using country level data on fuel-specific productivities but the requisite data do not appear to exist. Carbon intensities are computed as the ratio of carbon emissions to consumption for each fuel from the IEA data.

The initial values for capital, $K_{m,0}$, are chosen so that the capital stock is consistent with the model’s savings rate in the past. The initial level of labor productivity, $A_{m,0}$ then sets regional output equal to the data.

Figure 1 summarizes the regional part of the model calibration, with the model matching the high-level macro aggregates closely. Table 2 summarizes the targets used to determine the fuel- and region-specific parameters, with specific values reported in Tables A.1-A.4 in Appendix B.

Object	Target
<i>Aggregate production</i>	
Capital share, α_m	Matches labor share of GDP
Energy share parameter, ν_m	Matches energy expenditure share in GDP
Climate externality, ϕ_i	From Nordhaus (2010)
<i>Energy sector</i>	
Fuel-specific productivities, $\lambda_{m,i}$	Matches IEA fuel consumption
Carbon intensities, $g_{m,i}$	Matches IEA aggregate emissions
Research cost parameter, $\epsilon_{m,i}$	Matches IEA input costs
<i>Initial values</i>	
Capital stock, $K_{m,0}$	Savings rate \times Y_{t-1}
Labor productivity, $A_{0,m}$	Matches output
Reference cost of production, $\bar{p}_{m,i,0}$	Matches EIA-MER productivity growth

Table 2: Targets for region-specific calibration parameters and initial conditions. See Appendix B for further details.

Unmatched moments. The value of β is around 1 percent per year, consistent with standard values in the literature. This is chosen to match global supply and demand of conventional oil, and so its agreement with standard values represents an independent check on the model’s goodness of fit.

Two other unmatched moments are also useful quantitative validation of the model. First, the equilibrium oil price. The period 0 equilibrium oil price is \$528 per ton, or \$72 per barrel.¹⁹ The average during 2015-2019 was a little less than \$60. There is no parameter in the model targeting this moment – the real oil price is pinned down by a fixed global supply curve, and regional demands by the degree of substitutability with other fuels. That the global oil price is even close to the data is a success.

The other unmatched moment is the share of research in expenditures. The model imposes a

¹⁹Regional GDP is expressed in real 2012 US dollars, setting this as the numeraire for all other prices




Figure 1: Select region-specific calibration targets.

tight link between returns to scale and the share of energy expenditures that are accounted for by research: a fraction $1/(\eta - 1)$ of all energy expenditures go to fund research, around 3.5 percent in this calibration, or – given energy shares in GDP which of around 15 percent – approximately 0.5 percent of GDP. Data from the IEA’s World Energy Report suggests that this is perhaps a little high, as they imply a figure of around 0.17 percent of world GDP in 2018. However, these data only capture public and publicly-reported private R&D expenditures, so may be a lower bound on the true value. Broader measures of R&D expenditures provide a top-down alternative, with UN data estimating that around 2.1 percent of GDP is spent on research of all kinds. Under the assumption that research in energy is a similar share of expenditure to the economy as a whole, this should be comparable to the share of energy spending which goes to research, the 3.5 percent mentioned earlier. So the value in the model remains perhaps a little higher than ideal, but if energy is a more research-intensive sector than others (as seems plausible) this might not be a bad approximation. Given that this parameter is already constrained to meet data on returns to scale in research, it is at least reassuring that it is not wildly divergent from the data.

3.2 Laissez-faire solution

I start by solving the model under laissez-faire. That is, with no carbon taxes or green energy subsidies. To isolate the impact of endogenous technical change, I also solve a version of the model with fuel-

neutral exogenous technical change, where the input prices of all fuels decline exogenously at the rate of labor productivity growth (except conventional oil, whose price remains determined by global markets). The resulting paths for temperature and emissions are shown in Figure 2. The results from the exogenous technical change model are very closely aligned with other global integrated assessment models. Both Golosov et al. (2014) and Hassler et al. (2020) project a temperature increase of around 7 degrees Celsius – a catastrophic change in climate by any measure – and Golosov et al. (2014) annual emissions in 2200 of around 600 Gt CO₂. This is a useful check that despite the greater detail in terms of regions, fuel, and time frequency, the starting point for this framework is still consistent with the literature.

The addition of endogenous technical change to the model reduces emissions slightly over the model horizon, lowering the final temperature by a fraction of one degree. In other words, endogenous technical change per se is climate-neutral. Until around 2100, baseline endogenous technological change is increasingly dirty – global average emissions per unit of energy increase by almost one third. This is largely driven by increasing efficiency in coal, particularly in China. Innovation is a function of market size, and given that coal in China is by far the largest fuel market, innovations to technologies that increase the efficiency of coal usage have enormous potential scope. Without offsetting public policies, incentives to innovate in coal usage are thus very high, driving usage and emissions (see Figures A.1 and A.2 in Appendix B.2).

In the second half of the simulation, though, the composition of energy starts to shift more towards renewables (Figure A.2), eventually offsetting the increase in emissions intensity up to 2100. This is due to innovation in renewables, especially in Europe where the renewable energy share is a large enough encourage cost-saving innovation. Over time this increases the market for renewables, both in Europe but also elsewhere as improvements diffuse worldwide.

Figure A.1 also shows the different growth paths for the various regions. The most important sources of geographic heterogeneity are due to population growth and the climate externality. These typically act in opposite directions as warmer regions more often have higher population growth. But because the marginal cost of higher temperature is increasing, ongoing climate change acts as a larger drag on growth.²⁰ For example, sub-Saharan Africa has higher population growth than Russia, but a more severe productivity loss from climate change.

4 The domestic impact of climate mitigation policies

I now use the model to simulate outcomes under partially successful global climate policies. I divide this exercise into two stages, covered in this section and the next. In this section, I consider the domestic impact in the emissions-reducing regions: describing plausible emission reduction targets for Europe and China and discussing the economics behind the different policy combinations which achieve them (after all, the model allow for three different climate change mitigation policies). This analysis is important not only in its own right, but also because different ways of achieving emissions

²⁰A region-specific and often quadratic productivity-temperature relationship is common in the literature (Tol (2002), Nordhaus (2017)) and is consistent with evidence that productivity is a concave function of temperature (for example, Burke et al. (2015))




Figure 2: Global outcomes under laissez-faire.

targets can imply different technological spillovers to the rest of the world. Then, in Section 5, I analyze the mapping from these policies to global temperatures and global economic activity.

4.1 Domestic emissions targets

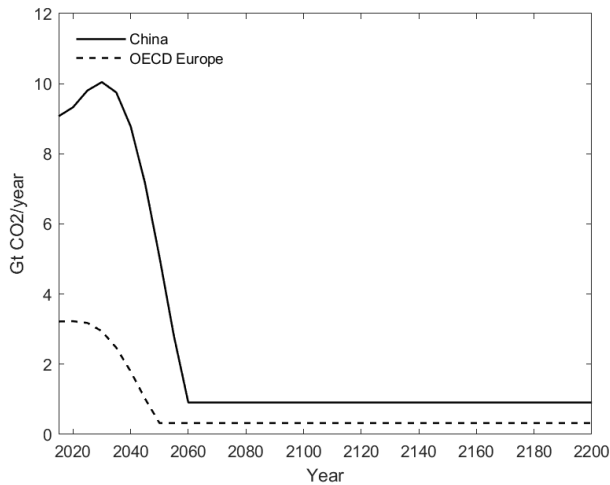
I consider a scenario where emissions which decline to 10 percent of their current levels for China and the European OECD model regions by 2060 and 2050 respectively, and remain at those levels thereafter. This is broadly consistent with announced policy plans.²¹ Most European OECD countries are members of the European Union, and so committed to its goal of net zero emissions by 2050. Those that are not EU members have similar goals of their own (e.g. the UK). Consistent with announced plans, I also require that Chinese emissions peak in 2030. These targets imply similar per capita emissions for the two regions; China’s lower per capita energy usage is almost exactly offset by a higher emissions intensity.

The specific target emissions paths are shown in Figure 3a. The mapping from regional emissions cuts to global temperatures is dependent on two effects. The first is a direct effect. China and OECD Europe account for around 40 percent of global emissions, so their reductions have a sizable impact on global temperatures. This direct effect is a result purely of emissions reductions in China and Europe and so is independent of the policies used to attain those reductions. Figure 3b shows this direct effect, recomputing global temperatures with lower emissions in China and OECD Europe but holding them fixed elsewhere. In line with the consensus view that regional policies are insufficient


²¹That emissions fall to 10 percent rather than exactly zero is for technical convenience, preventing marginal product of inputs from becoming infinite. Other interpretations are that this represents either a small overshoot of the target, or that thirty years from now carbon capture and storage is able to remove a small amount of CO2 from the atmosphere. Moreover, relative to laissez-faire, where emissions are increasing rapidly, the cuts are even larger – reducing emissions by over 96 percent in 2050.

to avoid catastrophic climate change, the effect is small – reducing temperatures by only around 1.5 degrees Celsius in 2200.

But there is also an important indirect channel. Technical change induced by policies in China and OECD Europe spill over to other regions, altering production possibility frontiers worldwide. A necessary first step in analyzing this channel is to discuss the policy mix in emissions-reducing regions. Emissions reductions can be achieved through different combinations of taxes and subsidies, and not all such policies have the same effect on the direction of future technological change. For example, a greater reliance on research subsidies will induce more technical change (and hence more international spillovers) than spending the same resources on production subsidies. How this indirect channel affects emissions in other regions – and thus global temperatures – is therefore a function not only of the magnitude of the emissions reduction in Europe and China, but also mix of policies used to do so.



(a) Emissions targets.



(b) Global temperatures.

Figure 3: Policy scenario. OECD Europe and China reduce cut emissions 90 percent by 2050 in line with announced plans.

4.2 Policies in emissions-reducing regions

I consider three candidate policy strategies to hit emissions targets. The first, a carbon tax only, where revenues are rebated to households lump sum. The other two are fiscally-neutral, where the proceeds of the carbon tax are funneled into either a green research subsidy, $\chi_{j,t}^m$, or a green energy production subsidy, $\delta_{j,t}^m$. For clarity, charts will label these experiments (1), (2), and (3) respectively. All three policy experiments of course deliver the same emissions paths.

Being able to back out the exact policies which deliver these emissions targets is one of the payoffs from assuming full depreciation of capital and log household utility. The computational strategy for inferring policies to meet emissions targets is discussed in more detail in Appendix D.2. However, the basic idea is to first fix the fractions of government expenditure on the two production subsidies and then search for the tax and subsidy policies that simultaneously hit the target path for emissions and balance the budget. Because the global oil price depends on global oil demand, policies in one region

affect contemporaneous outcomes everywhere else. Thus one needs to solve the full global general equilibrium model for each candidate policy in the search. With a fixed savings rate this can be done period-by-period, so the exercise reduces to repeatedly solving a four-dimensional root-finding problem, which can be done by multiple bisection.²² With a non-constant savings rate, the solution in each period would affect not only subsequent periods, but also previous periods – making this a highly non-linear 136-dimensional problem which is, for all practical purposes, impossible to solve.

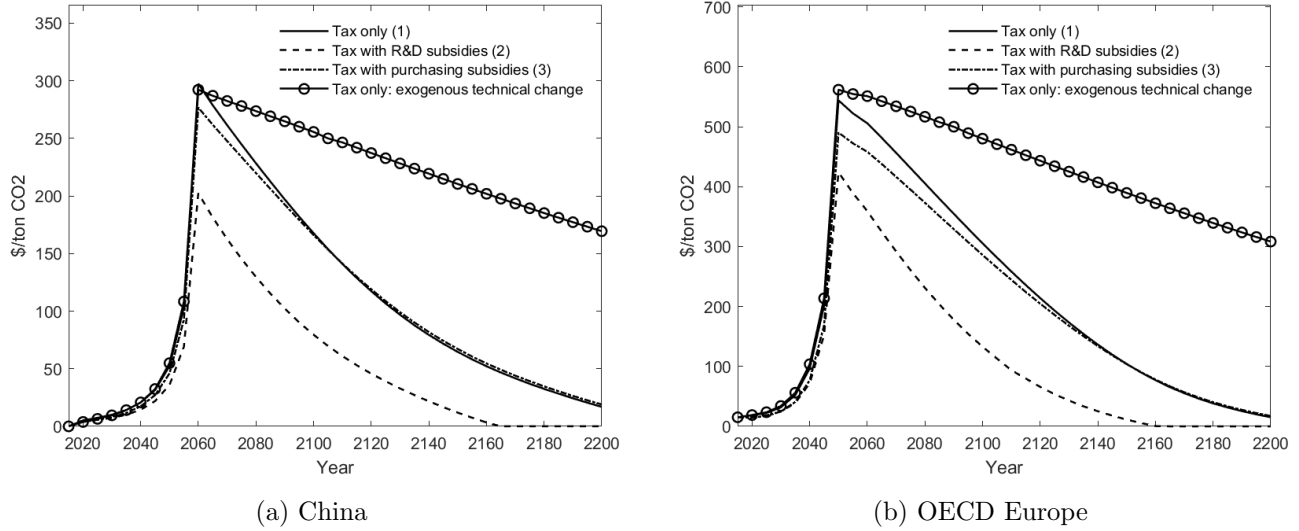


Figure 4: Carbon taxes required to meet emissions targets. Numbers in parenthesis refer to policy experiments, consistent across charts.

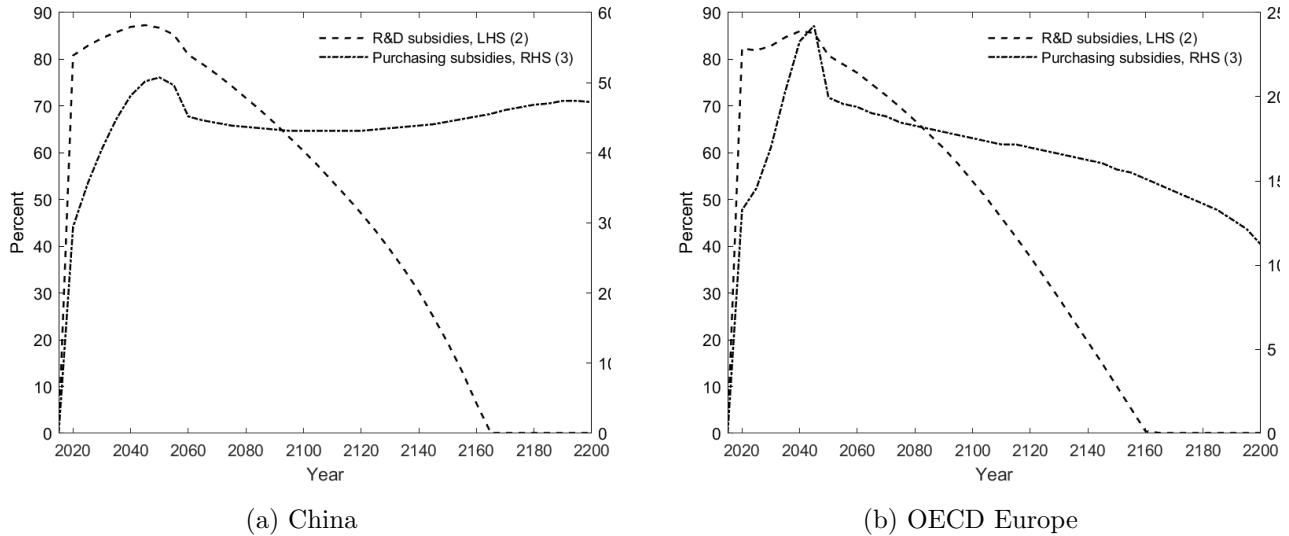


Figure 5: Balanced budget subsidies required to meet emissions targets. Numbers in parenthesis refer to policy experiments, consistent across charts.

Figures 4 and 5 show the policies required to hit the emissions targets in Figure 3a. Figure

²²The four equations are the emissions and budget balance targets in each region and the four controls the taxes and subsidies in each region.

A.4 in Appendix B.2 shows the resulting impact on household consumption. Looking first at the tax-only policy experiment – labeled (1) in figures – we see that taxes in both regions are initially low – remaining at or below \$100 per ton CO₂ in the first twenty or so years of the simulation – before increasing sharply as emissions decline. In addition, taxes in China are lower than in Europe. Endogenously directed technical change plays an important role in shaping the effectiveness of policies even without an explicit subsidy for green research. We can see this by comparing to the version of the model where technical change is exogenous and fuel-neutral. In this case, the carbon tax required to meet the same emissions target is much higher in the longer term, remaining at over \$200 per ton even in 2200. This difference occurs because with endogenous technical change, the carbon tax stimulates a shift towards green technologies in the medium and long-run, allowing taxes to be reduced much sooner.

Next, consider the case when tax revenues are spent on research subsidies, labeled (2) in figures. As shown in Figure 5, very high research subsidy rates are attainable, simply because carbon tax revenues are large and green energy research is small. The short term impact of subsidies is small; carbon taxes are very similar with and without subsidies prior to their peak. But over time, these subsidies shift technological change towards green fuels, reducing the peak carbon tax needed to meet regional emissions target by around one quarter to one third. In the longer run this also allows carbon taxes to be eliminated much more quickly – around 25 years sooner than under that tax-only case.

Research subsidies also mitigate the short-term welfare costs of climate change mitigation. Because they lower the cost of energy (and also because they permit lower future carbon taxes), household consumption falls by much less when carbon tax revenues are used to fund research subsidies, compared to a lump-sum rebate to households (Figure A.4 in Appendix B.2). In the long run, though, household consumption is higher for all climate mitigation policies simply because lower global temperatures reduce the impact of the climate externality.

Finally, we turn to the case when carbon tax revenues are spent on subsidies for purchases of green energy, labeled (3) in figures. In the short term, green research and purchasing subsidy expenditures are near-perfect substitutes, initially hitting the emissions targets with almost-identical carbon taxes.²³ But the dynamic responses are very different. I discuss this difference further in the next section when I compare the policies more formally. However, the intuition is simple. The immediate effects of production and R&D subsidies are very similar simply because they each redirect a similar amount of carbon tax revenue into consumers' pockets in the form of cheaper clean energy. But because R&D subsidies act through the technological component of input prices, their effect spills over to subsequent periods, raising the level of reference technology in future and lowering green energy prices in all future periods. In contrast, production subsidies provide little incentive for energy firms to innovate in green technology – green inputs are already cheap, so the marginal cost of savings from further research is much smaller, reducing inter-temporal spillovers. As a result, the future carbon tax can be barely reduced at all. In fact, the purchasing subsidy is so inefficient that households would prefer the higher consumption path under the tax-only policy (especially in China, see Figure A.4).

²³The large differences in the subsidy *rates* arises only because the eligible subsidy bases are of different magnitudes. Research is a fraction of the total cost of green energy and so similar expenditures on research and production subsidies entail very different subsidy rate.

4.3 The economics of domestic climate mitigation policies

Here, I discuss the economics behind the findings of the previous section, answering four important questions about the policy packages in the two emissions-reducing regions. These are: Why are subsidies ineffective in the short term? Why does the carbon tax rise so sharply in the medium term? Why are carbon taxes lower in China? And why are research subsidies more effective than production subsidies in the medium run?

To answer these questions, in Appendix C I develop a dynamic partial equilibrium model of the energy sector in a single country. This is a useful approximation for understanding the domestic impact of policies in the full model as it is very transparent, allowing for closed form solutions for the dynamic impulse responses to policy changes. These can be used to decompose of the effects of policy into its key channels, and highlight which parameters matter for the main conclusions. Hence, it is a useful vehicle for building economic intuition, which the the purpose for which i use it here.

Despite abstracting from important general equilibrium effects, the resulting framework still proves to be a good approximation to the dynamic effects of domestic policy. So in Section 4.4 I use this approximation to provide a complement to this intuition, building a full quantitative decomposition of the dynamic impact of policy in emissions-reducing regions.

4.3.1 Why are subsidies ineffective in the short term?

The near-term impact of subsidies is very small, barely reducing at all the requisite carbon tax in the first few decades (see Figure 4). To understand why, it is helpful to examine the contemporaneous marginal impact of a carbon tax and a production subsidy in the partial equilibrium setting. These are:²⁴

$$\frac{d \log M_t}{d \tau_t} = \underbrace{-\mathcal{E} \left(\frac{\eta}{\eta - \mathcal{E}} \right) \left(\sum_{i=1}^N (w_{i,t}^M - w_{i,t}) \frac{g_i}{\hat{p}_{i,t}} \right)}_{\text{Change in energy composition}} - \underbrace{F \left(1 - \frac{1}{\eta} \right) \left(\frac{\eta}{\eta - F} \right) \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}}}_{\text{Change in total energy consumption}} \quad (25)$$

$$\frac{d \log M_t}{d \delta_{j,t}} = \underbrace{-w_{j,t} \mathcal{E} \left(\frac{1}{1 - \delta_{j,t}} \right) \left(1 - \frac{1}{\eta} \right) \left(\frac{\eta}{\eta - \mathcal{E}} \right)}_{\text{Change in energy composition}} + \underbrace{w_{j,t} F \left(\frac{1}{1 - \delta_{j,t}} \right) \left(1 - \frac{1}{\eta} \right) \left(\frac{\eta}{\eta - F} \right)}_{\text{Change in total energy consumption}} \quad (26)$$

The only new notation here is are two weights: $w_{i,t}$ the weight of fuel i in expenditure and $w_{i,t}^M$ the weight of fuel i in emissions. The other key parameters are ones we have seen before: \mathcal{E} is the intra-fuel elasticity of substitution, F the elasticity of demand, η the curvature of the research cost function, and g_i the pollution intensity – the latter of which reflects the fact that the carbon tax relies on the carbon intensity of the fuel in question.

Equations (25) and (26) allow a decomposition of the impact of policy into two channels: the change

²⁴These formulas are linear approximations, exact at $\tau_t = 0$. In Appendix C I give general closed-form solutions for marginal impulse responses at all carbon tax rates. However, the formula are clearest when $\tau_t = 0$, i.e. for the first dollar of carbon tax, or a subsidy without a carbon tax. I thus use this example for exposition. I also focus on the production subsidy as representative of the two subsidies here as the contemporaneous effect of a research subsidy is very similar. Given that we are only thinking about the domestic impact of policies, I also suppress the region-specific index m .

in the composition of energy and the change in the total energy consumption. These two channels have an important economic interpretation – the first is driven by relative prices of inputs and the second by their average price, akin to the familiar substitution and income effects in a consumer choice problem. Together these channels determine the overall response of emissions.

This decomposition immediately makes clear the most obvious difference between the two policies; that the two channels go the same way for a carbon tax (both negative) but oppose each other for the subsidy. This is because both policies increase the relative price of dirty fuels, causing a substitution into cleaner fuels, but only the carbon tax raises the average price, driving a reduction in total energy consumption. In contrast, the green energy subsidy lowers the cost of energy overall, increasing energy consumption. This means that a green energy subsidy is a less effective tool than a carbon tax in meeting an emissions target. Indeed, if the elasticity of demand F is high enough (or the elasticity of substitution \mathcal{E} low enough) green energy subsidies can boost emissions, necessitating a higher carbon tax to hit the emissions target.

There is a second important factor that these formulas highlight: that effect of a subsidy is proportionate to the size of the green energy sector, $w_{j,t}$. In most regions this is currently low, and while subsequent green technical change may cause this to grow, this seriously constrains the effectiveness of the green energy subsidy in the model in the short term (which is calibrated to hit national energy shares). Indeed, this explains why the subsidies are ineffective in the model even though the choice of parameters means that the composition channel is large and the quantity effect small (recall: the elasticity of substitution $\mathcal{E} = 2$ is much larger than the elasticity of demand $F = 0.35$).

Two other points are worth noting from these formulas. First, that the induced innovation multiplier for each policy is the same: $\eta/(\eta - \mathcal{E})$ for the composition channel, and $\eta/(\eta - F)$ for the quantity channel.²⁵ This is because the indirect incentives for innovation are similar. Innovation is induced by an increase in the cost share of a particular fuel (equivalent to the “market size” in Acemoglu et al. (2016) and Fried (2018)); the more that energy-producing firms spend on a particular input, the more incentive they have to save costs by innovating in the use of that input. This multiplier is a function of returns to scale in research; at lower returns to scale (i.e. larger η) this multiplier decreases and limits to 1 as $\eta \rightarrow \infty$.

Second, that the quantity channel explains why subsidies, even if ineffective in reducing emissions, can mitigate the welfare costs of climate mitigation policies in the short term. By lowering consumers’ energy costs, subsidies allow higher energy consumption (seen in the positive quantity effect). This can partly offset the reduction in energy consumption due to higher carbon taxes (seen in the negative quantity effect), and explains why subsidies can raise household consumption (see Figure A.4 in Appendix B.2).

4.3.2 Why does the carbon tax rise so sharply in the medium term?

A common feature of the policies required in both China and Europe is the very sharp increase the carbon tax rates required to reduce emissions in the medium term. Of course, part of this is because

²⁵The $(\eta - 1)/\eta = 1 - 1/\eta$ terms adjust for the fraction of the final energy price which is due to research, which is always $1/\eta$.

the decline in emissions is quite steep. But this is not the whole story; the increase in taxes is disproportionately larger than the decline in emissions. Peak to trough, emissions in both regions fall by around a factor of ten. Over the same period, taxes increase by a factor of thirty. This is the result of a fairly simple static effect: that the marginal effectiveness of the carbon tax decreases as it gets higher. And so ever higher taxes are needed to meet ambitious emission reduction targets.

Three mutually reinforcing mechanisms drive this. First, because the carbon tax deters pollution, the base for the tax shrinks. Thus, reducing emissions by a given volume requires ever higher percentage increases in post-tax input costs. Second, because taxes themselves drive up post-tax input costs, a given percentage increase in input costs requires a larger increase in carbon taxes. Third, at high carbon taxes, input costs become almost proportional to the tax for all dirty fuels. As a result, the carbon tax no longer discriminates between different dirty fuel types, and suppressing the composition effect within dirty fuels.

Figure 6 illustrates this point, computing the static (i.e. within-period) response of emissions and input usage to a carbon tax in the initial period in China. The marginal effectiveness of the first dollar of a carbon tax is highly effective, but this declines sharply, particularly once the tax is over \$100 per ton (Figure 6a). A carbon tax of around \$300 is required to deliver reductions in the order of 90 percent, i.e. roughly the peak of the policy path in Figure 4a.

The reason that the carbon tax is so effective at low levels is that it discriminates very sharply between different fuel types. In particular, even small carbon taxes cause a very notable drop in coal usage (Figure 6b). To understand why this happens, equation (25) is once more useful, in particular the composition channel which relies on two factors. The first is the weight of the fuel in emissions $w_{i,t}^M$. The inclusion of this is purely mechanical, as the impact on emission of a given percentage change in fuel usage scales with the share of emissions. The second is the implicit subsidy to the fuel under laissez-faire, $g_i/\hat{p}_{i,t}$. Literally, $g_i/\hat{p}_{i,t}$ is the number of tonnes of CO2 that a single dollar can put into the atmosphere; if CO2 were the intentional output of energy it would be the productivity of a given fuel. But because the social benefit of CO2 is negative, we can also interpret $g_i/\hat{p}_{i,t}$ as the implicit subsidy on pollution (up to addition of a constant). That this shows up in the marginal response of fuel usage is due to the Pigouvian nature of the carbon tax; by taxing the source of pollution directly it offsets the implicit subsidy on pollution. At around four times the next-worst fuel, coal has by far and away the highest implicit subsidy $g_i/\hat{p}_{i,t}$. Put simply, coal is dirty (g_i is high) and cheap ($\hat{p}_{i,t}$ is low). Thus, the marginal effect of low carbon taxes on coal (and hence emissions) is very large. Table A.5 in Appendix C provides a full decomposition of the first dollar of the carbon tax into the composition and quantity channels for each fuel, laying out this arithmetic explicitly.

However, at higher taxes, the carbon tax fails to discriminate between the different dirty fuels. The post-tax price is mostly tax, so a one dollar increase in the tax produces the same percent increase in the price of all fuels. As a result, the usage of all dirty fuels decline in constant proportion at high tax rate (see again Figure 6b). This declining marginal effectiveness of the tax is why it has to rise so sharply to squeeze out the last bit of fossil fuels being used.

The steeply declining effectiveness of the carbon tax also has implications for how we view policy recommendations when there is model uncertainty. The near-zero elasticity of emissions with respect

to the tax means that if the model is slightly mis-specified (i.e. if the line in Figure 6a is slightly higher or lower) then the carbon tax required to hit a given emissions target will vary considerably. In contrast, there is less uncertainty about the tax required to meet a less ambitious emissions target. This means that one should not fixate on the *peak* of the carbon tax shown in Figure 4, but instead the low and slowly growing level in the first decade or so. This basic intuition not special to this model, but applies more broadly to other models with carbon taxes. Indeed, it helps to understand why there is enormous variation in the optimal carbon tax in the literature, as small changes in emission-reduction targets imply huge differences in the requisite tax.²⁶




Figure 6: Static responses to a carbon tax, China.

4.3.3 Why are carbon taxes lower in China?

Figure 4 shows that the requisite emission-reducing taxes in China are substantially lower than in OECD Europe. Once more, equation (25) can help us understand why. The intensive usage of coal in China, accounting for around two thirds of all energy, generates a correlation between the emissions share $w_{i,t}$ and the implicit subsidy $g_i/\hat{p}_{i,t}$. As a result, the composition channel in China is much stronger than in Europe, and so the same emissions reductions can be accomplished with smaller lower tax increases. More intuitively, China has a much larger carbon tax base and so can generate the same reduction in emissions with much smaller increases in the tax itself.

4.3.4 Why are research subsidies more effective than production subsidies?

Another notable feature of Figure 4 is that spending tax revenue on research subsidies is much more effective than spending on production subsidies, in the sense that it permits a larger and sooner reduction in carbon taxes. This seems like a potentially policy-relevant finding, so here I walk through the drivers of this a little more. The key intuition is that the contemporaneous effects of research and

²⁶These range from just \$6 in Nordhaus (2014b) to \$85 in Stern and Stern (2007) and even higher (e.g. Council (2013)). See the survey in Van den Bergh and Botzen (2015) for further details.

production subsidies are near-identical. But because a production subsidy introduces a wedge between productivity and firms' input costs, it does not incentivize research to the same extent. This causes a difference in the dynamic responses, as future research gains build on past research effort. Over time, this difference can become quite large, resulting in the differences seen in Figure 4.

Again, a partial equilibrium approximation can make this clearer. As I show in Appendix C, the response of aggregate energy usage and emissions to the two types of subsidy are in fixed proportion:²⁷

$$\frac{d \log E_t}{d \delta_{j,t}} = (\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log E_t}{d \chi_{j,t}} \quad (27)$$

$$\frac{d \log M_t}{d \delta_{j,t}} = (\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log M_t}{d \chi_{j,t}} \quad (28)$$

This means that the two subsidy policies are substitutes at rate $(\eta - 1)(1 - \chi_{j,t})/(1 - \delta_{j,t})$, in the following sense: a change in policy which simultaneously reduces (increases) the purchasing subsidy by one percentage point and increases (reduces) in the research subsidy by $(\eta - 1)(1 - \chi_{j,t})/(1 - \delta_{j,t})$ percentage points has no effect on contemporaneous aggregate emissions and energy.

Of course, such a policy change might be expensive. To understand the fiscal impact, we can compute the marginal effect on the government surplus s_t if they increase research subsidies and reduce purchasing subsidies in the proportions which leave aggregate variables unchanged.

$$(\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log s_t}{d \chi_{j,t}} - \frac{d \log s_t}{d \delta_{j,t}} = \frac{\chi_{j,t} \eta}{\chi_{j,t}(1 - \delta_{j,t} \eta) + \delta_{j,t}(\eta - 1)} - \frac{1}{1 - \delta_{j,t}} \quad (29)$$

This is typically very small, and exactly zero when $\chi_{j,t} = \delta_{j,t}$. So an extra dollar of public expenditure on either policy will have very similar effects, and identical ones when research and production subsidies are equal. In this sense, research and production subsidies are perfect policy substitutes. Switching between the two at the appropriate rate will leave not only environmental and economic outcomes unchanged but also the fiscal position as well.

To understand the intuition for this result, note first that the marginal impact of a higher subsidy on post-subsidy costs is nonlinear. This is true for any subsidy: a one percentage point increase in the subsidy rate reduces relative costs by one percent if applied to an initial rate of zero, but by 2 percent if subsidies already over half of costs. And so a one percent increase in both subsidy rates has the same impact on the costs of production and research only if they are equal, i.e. $\delta_{j,t} = \chi_{j,t}$. But this only really replaces one question with another: conditional on initial subsidy rates being the same, why do marginal changes in the two subsidies have the same real impact? The answer is that two opposing effects distinguish a dollar spent on research subsidies from a dollar spent on production subsidies. These effects happen to offset identically in the model. The first effect is that a dollar of research subsidy goes much further than a dollar of production subsidy. Whereas production subsidies apply to all energy spending, research subsidies only apply to the component of energy costs accounted for by research. Thus, the same dollar can fund a much bigger increase in the research subsidy rate. Because the research share of energy costs is $1/(\eta - 1)$, a dollar of subsidy can fund a reduction in

²⁷Again, I focus on the case where $\tau_t = 0$, for transparency

research costs that is $\eta - 1$ times as large as if spent on production subsidies. The second effect is that a one percent improvement in research efficiency does not lower production costs by one percent, but instead by the returns to scale in research, $1/(\eta - 1)$. Thus, relative to a production subsidy, a dollar of research subsidy can fund a reduction in research costs that is $\eta - 1$ times as large but has an effect $\eta - 1$ times smaller. The effects offset identically. That the research share and returns to research are correlated makes sense: higher returns mean that firms want to spend a higher fraction of costs of research. That they are a similar size is not empirically unreasonable, as discussed in Section 3.1.

There is also a dynamic effect from changing subsidy policy. A sufficient statistic for the future fuel mix is the relative input costs of clean and dirty fuels, $p_{j,t}/p_{i,t}$, which determines future relative input prices. A change in subsidy policies leaving E_t, M_t unchanged will alter this as follows:

$$\begin{aligned} (\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log p_{j,t}/p_{i,t}}{d \chi_{j,t}} - \frac{d \log p_{j,t}/p_{i,t}}{d \delta_{j,t}} &= (\eta - 1) \left(\frac{1}{1 - \delta_{j,t}} \right) \left(\frac{-1}{\eta - \mathcal{E}} \right) - \left(\frac{1}{1 - \delta_{j,t}} \right) \left(\frac{1 - \mathcal{E}}{\eta - \mathcal{E}} \right) \\ &= \frac{-1}{1 - \delta_{j,t}} < 0 \end{aligned}$$

So the research subsidy lowers the future cost of green fuels relative to dirty ones by more than a purchasing subsidy. This happens despite identical quantity responses because the production subsidy acts as a wedge between the input cost and production (i.e. technological) cost of the green fuel. And although both policies induce some green innovation, the effect of the research subsidy is entirely through this channel. These differences compound over time, affecting the whole future path of technological progress.²⁸

Overall, this analysis explains why research subsidies are much more effective than production subsidies. The contemporaneous effects of the two policies is similar, and identical when research and production are subsidized at the same rate. However, this is dwarfed by the dynamic effect, which compounds over time to make research subsidies wildly more effective.²⁹

4.4 Decomposing the dynamic effect of policies: domestic impact

To understand more precisely the exact dynamic effect of policies, Figure 7 and Figure A.3 (in Appendix B.2) use the impulse responses of the partial equilibrium setting to provide a quantitative policy decomposition of the change in emissions and energy usage until 2100 in the most successful policy experiment discussed above; that where carbon tax revenues are used to fund research subsidies

²⁸Indeed, I show in Appendix C that if intra-fuel spillovers are weak, one fuel can attain an insurmountable technological lead, with a self-reinforcing circle of ever-increasing energy share stimulating ever-higher research driving down costs of production more and more. This is in fact exactly the intuition behind the mechanism in Acemoglu et al. (2016), who show that temporary climate mitigation policies can lead to permanent changes in fuel shares in the long run. With sufficiently large technological spillovers, advances in one fuel drive down costs of production in the others, leading to common movements in production costs and preventing runaway progress in one fuel type. In a special case of the partial equilibrium model, it is possible to provide a specific lower bound on the degree of spillovers, expressed in terms of the eigenvalues of the spillover matrix. See Appendix C.4 for details.

²⁹One obvious question is: why not also show a simulation where $\chi_{j,t} = \delta_{j,t}$ always? After all, this would produce the largest contemporaneous effect on emissions. The answer is that because the dynamic effect is so much larger than any contemporaneous differences between the policies, that only differences in $\chi_{j,t}$ matter and this exercise ends up almost exactly halfway between experiments (2) and (3). Including this as a fourth option throughout would add more complexity than insight.

(labeled (2) in Figures).

To calculate the policy decomposition, in each period, I compute the partial equilibrium impulse response due to the policy changes in that period relative to *laissez-faire*, using a finite difference method to build up nonlinear impulse responses from the linear approximations. The sum of these period-policy specific impulse responses defines the total partial equilibrium response.³⁰ The results of this exercise are very close to the general equilibrium outcomes (see “approximation error” in Figures 7 and A.3).³¹ Within each policy, I distinguish between the three channels discussed above: the static composition and quantity effects, and the (dynamic) induced technical change. These are important because they changes in three fundamental economic concepts, respectively: the relative prices of fuels, the average price of fuels, and incentives for research. Moreover, these channels are the primary avenues through which key parameters to impact policy conclusions: the elasticity of substitution \mathcal{E} principally drives the composition effect, the elasticity of demand F impacts the quantity effect, and the convexity of the research cost function η determines induced technical change.

In both China and OECD Europe, the contemporaneous impact on emissions of the carbon tax far outweighs that of the research subsidy in all periods. This is just the mechanisms discussed in Section 4.3.1 at work; the contemporaneous impact of the carbon tax is large because it discriminates sharply between fuels, and that of the subsidy small because the expenditure weight of green energy is small. This is reflected in Figure 7 by the large static composition effect of the carbon tax – the quantity effect is almost zero.

The dynamic effects of the tax and subsidy policies, however, are much more similar. Because the research subsidy directly targets the technological component of input costs, it induces a technological response similar to that of the carbon tax, despite a much smaller contemporaneous impact on quantities. As a result, the emissions impact of technological change induced by research subsidies and carbon taxes is of approximately the same size, evidenced by the similar sizes of the green and brown bars in Figure 7. In other words, spending carbon tax revenues on research subsidies roughly doubles the induced emissions effect versus the carbon tax alone.

The impact on energy consumption (the partial equilibrium proxy for changes in household consumption) is a very different, with research subsidies and carbon taxes operating in different directions, and the research subsidy explaining almost all the long-run improvement in household consumption in the medium run (see Figure A.3). The difference versus the emissions responses reflects their differing importance of the energy composition and quantity channels. Whereas the composition channel is crucial for emissions, it is irrelevant for total energy use – literally, pure composition changes are by definition orthogonal to quantity changes. Hence the absence of the purple bars in Figure A.3. Because the carbon tax raises the price of energy overall, it reduces total energy consumption both contemporaneously and – by lowering the overall level of energy technology – in future. In contrast,

³⁰Of course, one could do a similar experiment using the full general equilibrium model, building up the impulse responses by re-solving the model with each policy turned on and off in each region and period. But even with just two regions, two active policies, and almost forty periods, this quickly becomes computationally burdensome. In contrast, the partial equilibrium approach solves almost instantly on a modern laptop with only a small approximation error. Appendix C explains the policy decomposition method in more detail.

³¹This error term includes both within-region general equilibrium effect via aggregate demand, international pecuniary spillovers via the global oil price, and the climate externality.

the research subsidy permanently lowers the price of energy overall, acting against the impact of the carbon tax increasing energy consumption, especially in the long run.

In summary, this exercise makes it clear why the carbon tax plus research subsidy program is so effective. The carbon tax acts to quickly reduce emissions in the short term and induce green technical in future, but comes at the cost of making energy more expensive for households. The research subsidy is less effective at reducing emissions in the short-term, but has a long-run impact equal to that of the carbon tax. The short-term value of the research subsidy comes from cushioning welfare cost of the carbon tax by lowering the cost of energy.




Figure 7: Policy decomposition of changes in emissions relative to laissez-faire for a tax with R&D subsidies, experiment (2). Using analytical impulse responses from a partial equilibrium approximation to the model, the impact of each policy is divided into three channels: static composition and quantity effects, and the dynamic induced technical change effect.

5 The impact of international technological spillovers


I now turn to the global impact of the regional emissions-reducing policies.

5.1 Global temperature response to regional policies

First, in Figure 8a, I report the global temperature paths associated with the different policy combinations discussed in the previous section. The most obvious point to note is that no matter how China and Europe achieve their emissions reductions, the effect of international technological spillovers is very large, especially in the long run. Even the worst case – when emissions are reduced through a carbon tax alone – temperatures in 2200 are reduced by nearly another 2 degrees beyond that due to the reduction in emissions in Europe and China alone.

A crucial mechanism here common to all policies is the elimination of Chinese demand as a driver of greater efficiency in coal usage. As a result, in the model China drives global innovation in coal technologies, with coal input costs on average around 10 percent below the next cheapest region (see Figure 8b). In the laissez-faire scenario, this innovation is a negative externality, as it diffuses to other regions and reduces the cost of pollution worldwide. Because the carbon tax common to all three policy experiments reduces demand for coal, incentives to innovate are greatly reduced. By 2050, China is no longer driving global productivity improvements in coal (the line in Figure 8b crosses the x-axis) and this negative technological externality is eliminated.

In the next section I give a more precise accounting of the role of coal innovation in driving of these spillovers. But given the importance of this channel, it is natural to wonder immediately exactly what coal-specific technologies might potentially be dissuaded by the various policy packages. The model is too general to make specific predictions on this point. And in any case, forecasting new technologies many decades is highly speculative. That said, one can easily imagine a number of possible developments which might reduce the cost of producing energy from coal. For instance, automation might allow faster and cheaper mining; or improvements in communication or transport reduce shipping and storage costs; or improved materials science reduce wear and maintenance costs for coal power plants. Perhaps more convincing than speculating on particular advances, though, is the more general point. That is, that the Chinese coal market is simply enormous, with expenditures equivalent to approximately 2 percent of GDP. This means that even small efficiency improvements can deliver potentially enormous cost savings, in the real world as much as in the model, creating correspondingly large incentives for innovation. Given this, it seems not just reasonable to expect that this will drive innovation and cost reduction in this sector for years to come, but perhaps even unreasonable to expect that it would not.



(a) Global temperatures relative to pre-industrial average

(b) Coal input price in China, percent difference relative to cheapest other region

Figure 8: Effect of changes in the policy mix in Europe and China.

Returning to Figure 8a, there is some non-trivial variation in global outcomes across the different

policy packages. Funneling the carbon tax back into research subsidies is effective in reducing temperatures yet further by around a further 1 degree in 2200. This is due to two opposing factors. On the one hand, green research subsidies somewhat surprisingly reduce the price of coal in the long run (see Figure 8b). This occurs because they lower the overall price of energy in the long run and so increase the market for all fuel types. This eventually results in lower coal input costs than a carbon tax which spill over to other regions, lowering the cost of pollution. However, this is offset by a more obvious direct effect – more green research lowers the cost of green energy which spills over internationally. This latter effect dominates, lowering global temperatures.

Although production subsidies are not a terribly useful part of the policy mix from the perspective of meeting domestic emissions targets, they can also have positive long run international spillovers. Domestically, the substitution into green energy that they induce is partly offset by an increase overall energy consumption. But internationally, production subsidies can contribute to climate mitigation via international spillovers, lowering temperature by a little less than one half of one degree in 2200. This is because the offsetting domestic factor (lower prices and higher consumption for energy) is absent in the international transmission. Other countries see only the extra research induced by production subsidies. As this is greener than the global average, this makes technological progress worldwide greener.

5.2 Decomposing the dynamic effect of policies: international spillovers

To make clearer the relative importance of technological spillovers from clean and dirty fuels, I focus once more on experiment (2), explicitly decomposing the impact of spillovers into different channels – in this case the impact of fuel-specific technological progress. This is a more quantitatively specific analysis of the mechanisms in the preceding discussion and is the international complement to the domestic analysis in Section 4.4. The key evidence in this discussion is Figure 9, which decomposes emissions in two other major emitters into the contributions from advances in different fuels. Here I focus on India and the United States as they are the two next largest-emitting regions and are broadly representative of emerging and advanced economies respectively.

Intuitively, Figure 9 traces through the impact of the changing productivity of each fuel on each region by repeatedly re-solving the model with exogenous input prices, gradually shifting input prices fuel-by-fuel from those which arise with endogenous technical change under *laissez-faire* to those in experiment (2), a carbon tax which funds R&D subsidies. More formally, I construct a sequence of fuel-region-time input price paths $\{\mathbf{p}_{m,t}^1\}_{t=1,\dots,T}^{m=1,\dots,M-1}, \dots, \{\mathbf{p}_{m,t}^N\}_{t=1,\dots,T}^{m=1,\dots,M-1}$ where: the first element of this sequence is the input prices attained under *laissez-faire*; the last element of this sequence the input prices which attained in policy experiment (2); and the intervening sequence is monotone in each element. I then solve a version of the model with exogenous technical change for each of these price paths, changing fuels one-by-one and recording the incremental impact on emissions in each region. The sum of these incremental impacts then defines the contributions in Figure 9.³²

The resulting cumulative change in emissions is a very good approximation of the change in emis-

³²I use a sequence with 100 steps. With six improvable fuels this means solving the model 600 times. The simplicity of the model is once again essential to making this a feasible exercise.

sions in the full endogenous technical change model for the non-active regions (i.e. not China or OECD Europe); the Approximation Error bar in Figure 9 is almost invisible.³³ And while the fuel-by-fuel decomposition also captures some second round effects, such as incentives for innovation in renewables in the non-active due to a larger domestic market, the primitive source of the shock is a change in foreign policies. It thus reflects the impact of changing input prices in India and the United States in response to policies overseas. That is, international technological spillovers.

Figure 9 makes precise the arguments of the preceding section. The the bulk of the reduction in emissions in non-active regions is due to changes in the input cost of just two fuels: coal and renewables. Even though the changes in input prices go in opposite directions (up for coal, down for renewables) their contribution to emissions have the same sign because coal is a polluting fuel type and renewable energy is not. Quantitatively, the deterrence of improvements in coal technologies are typically a little more important, accounting for around two thirds of the decline in emissions. The split in other regions (not shown) is very similar.




Figure 9: Decomposition of changes in emissions relative to laissez-faire for two major inactive regions. Tax with R&D subsidies, experiment (2). Effects are due to fuel-specific technological change spillovers from China and OECD Europe, attributed to individual fuels by exogenously shifting input costs for each fuel in each region (see text for details).

Overall, the headline results suggest that international technological spillovers can be a very effective substitute for international policy cooperation, in the best case delivering long-run temperatures below 3 degrees above pre-industrial norms, versus over 7 degrees in the baseline. Perhaps the most interesting aspect of this is the decomposition of this effect. Of the roughly 4.5 degree reduction in long-run temperatures, around 3 degrees is due to international technological spillovers. But of this,

³³Specifically, this will get wrong the global general equilibrium effects coming from changes in global temperature and oil prices. But given that the bulk of the effect on non-active regions is from spillovers – which they take as exogenous – it is a good approximation for the non-active regions.

incentives for new green technology play a secondary role, accounting for only around 1 degree Celsius in 2200. Instead, the most important channel by which international technological spillovers amplify the effectiveness of policy is by dissuading productivity improvements in dirty fuels, especially coal.

5.3 Alternative parameterizations of international technological diffusion

To help understand the importance of international technological diffusion in these results, I consider five alternate calibrations of international technological spillovers. Figure 10a presents the resulting change in temperatures in these experiments, expressed relative to laissez-faire to isolate the impact of policies (since changing parameters also changes the laissez-faire baseline).

The first alternative calibration is one which shuts down entirely the international diffusion of technology, setting $\zeta_m = 0$ for all regions. This is a useful benchmark, because it addresses a reasonable critique of the explanation so far. By changing emissions and altering international demand for oil, policies in emissions-reducing regions affect the marginal cost of energy and the impact of the climate externality. In principal, these general equilibrium effects other than international technological spillovers could be driving the results in Figure 8a. Figure 10a shows that this is not the case, with temperatures in a model with zero spillovers – and thus includes the general equilibrium channel – almost identical to the direct effect of reducing emissions in China and Europe, lowering temperatures in 2200 by less than 0.5 degrees Celsius. The difference between the two is principally the climate damages. That is, the endogenous feedback of temperature to output which lowers future productivity and emissions, pulling down on temperatures.

The second alternative calibration is one with lower spillovers from endogenous technical change. Given that international spillovers seem so powerful, it seems wise to try a version with smaller yet still plausible spillovers. Eaton and Kortum (1999) report a second value for the speed of international spillovers – the unconditional rate of adoption on foreign technologies, which they estimate to be 21 years. This gives a value of $\zeta_m = 0.24$ for all regions. The resulting impact on long-run temperatures is still very large – a decline of almost 4 degrees Celsius or around double the direct effect. This exercise also tells us something about the marginal impact of changes to spillovers. This value of ζ_m is about halfway between the headline and no-spillovers cases, yet the long-run temperature impact is around two-thirds of the way to the headline result. Thus, the relationship between spillovers and long-run temperature is likely concave, with small spillovers have larger *marginal* effects on long-run temperature. This should give us extra confidence in the results, as it implies that spillovers have to be very small indeed to invalidate the broader conclusion that international technological diffusion significantly lowers long-run temperatures.

The third alternate calibration draws on a line of literature which emphasizes the value of trade links – specifically imports – in facilitating the transfer of technology. Examples include Caselli and Coleman (2001), Sampson (2016), Bloom et al. (2016), and Buera and Oberfield (2020), the last of whom develop a model where the flow of ideas from one country to another is a function of flow of

imports in the same direction. In order to capture this idea, I replace equation (19) with the following:

$$\bar{p}_{m,i,t}^* = \sum_{k=1}^M w_{m,k}^I \min \left(\bar{p}_{m,i,t}^d, \bar{p}_{m,k,t}^d \right) \quad \forall i, m \quad (30)$$

Where $w_{m,k}^I$ is the share of imports in region k coming from region m . The interpretation is that for each fuel type, every country has the chance to learn best practice from its trading partners, and that these opportunities are in proportion to the import share for each trading partner. This addresses the concern that equation (19) is permitting too simple a transfer of the very best technologies from the rest of the world. Here, the technological frontier is not simply the most productive global technologies, but instead a weighted combination of them. I compute the import shares using UN COMTRADE data for the initial model period, 2015-2019.

Allowing technological diffusion to follow imports does not make a huge difference to the headline results. Long-run temperatures still decline by almost 4 degrees Celsius. The reason is fairly simple – China and OECD Europe together account for around 50 percent of imports into other regions. So although changing the speed of diffusion to reflect trade links will limit policy spillovers to some extent, policies in these two regions continue to have a substantial effect on the path of technology elsewhere, simply because these regions are major exporters.

The last two experiments involve increasing the rate of technological diffusion so that all regions fully catch up to the frontier with a one period lag. In one experiment diffusion is fast for all fuel types, and in the other it applies just to green fuels. We can think of these as a rough way of trying to answer a natural question: can the positive impact of international technological spillovers be exploited somehow, perhaps through programs to spread new knowledge that can improve energy efficiency? These two experiments explore the outer bounds of what might be possible with such programs. In both cases the effect is relatively small, again reflecting the declining marginal effectiveness of ever-faster technological diffusion.

The speed of international technological diffusion also affects how the burden of climate mitigation policies are shared. Although faster diffusion amplifies the global benefits from regional climate mitigation policies, it also makes meeting targets harder for the regions which chose to pursue them. Because technological progress flows not only from Europe and China to the world, but also from the world back to Europe and China. Because the emissions-reducing regions are implementing policies which hinder dirty technological development, they end up behind the global frontier in dirty energy production. And so knowledge about innovations in these fuels flows in to these regions while innovations in green technology flow out. As a result, faster technological diffusion reduces the cost of fossil fuel inputs in Europe and China, necessitating higher carbon taxes. As an example, Figure 10b shows the resulting carbon taxes in Europe for a subset of these experiments, showing that high spillovers for all technologies would necessitate an increase of around \$50/ton CO2 in the short term, and long-run cuts in carbon taxes are delayed by around two decades. In contrast, in the absence of international spillovers, European countries would be able to meet their emissions targets while cutting carbon taxes around a decade earlier than in the headline case. Of course, this comes at considerable cost to the global environment, as we saw earlier.

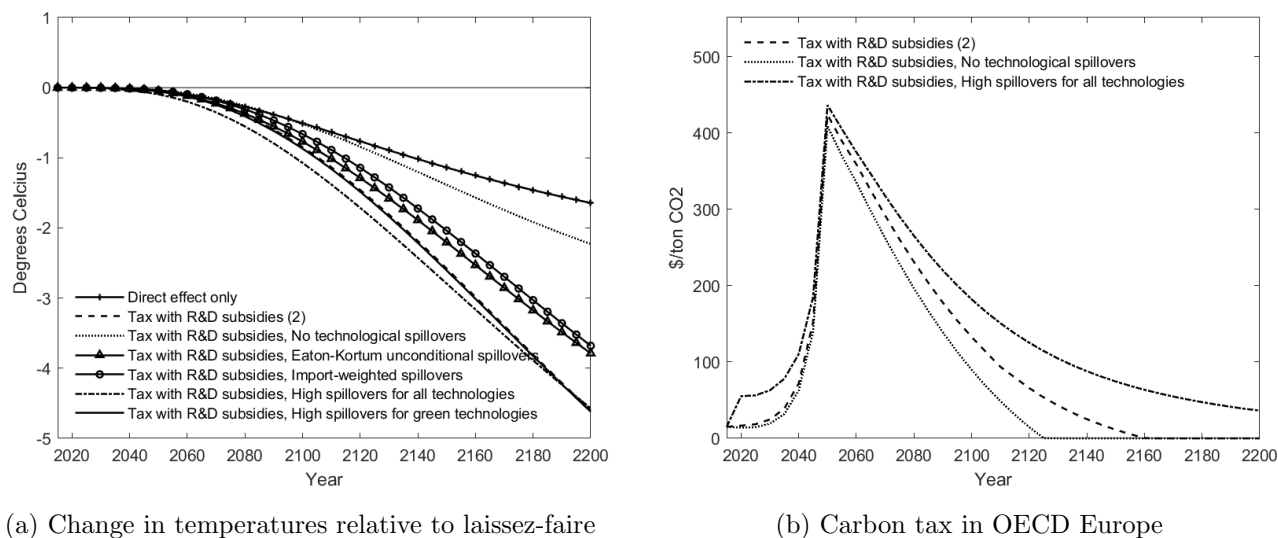


Figure 10: Effect of speed of technological diffusion.

6 Robustness

In order to be able to invert the model to back out policies, I make some fairly strong modeling assumptions. Perhaps the most stark are those about the production function for research and the full capital depreciation and log preferences of the household. I now check the robustness of my main results to these assumptions and show that they do not substantively impact the main results.

6.1 Alternate forms of endogenous technical change

An important set of modeling assumptions that I made were on the research sector. I assume that research is an unremunerated factor of production (i.e. a pure cost, like an imported intermediate) and has a simple cost structure with constant returns to scale. The first of these guarantees that households derive no income from payments to the energy sector, which is an important part of guaranteeing a constant savings rate. The second is purely for simplicity.

To check that these do not drive my results, I recalibrate my model to match as closely as possible the deep parameters of two cutting-edge global integrated assessment models, Fried (2018) and Acemoglu et al. (2016). These both feature more sophisticated models of how innovation is produced, using households' labor as an input. I then replicate the baseline and policy experiments reported in both papers and show that they correspond closely to those in my model. From this, I conclude that the simplifications that I make do not quantitatively affect the policy responses of the model. As such, the main conclusions for policy are robust to alternate specifications for endogenous technological change.

6.1.1 Comparison to Fried (2018)

Fried (2018) studies the interaction of endogenous directed technical change and a carbon tax. Hers is a global (i.e. one-region) general equilibrium model with two fuel types: fossil fuel and green energy.

There is endogenous technical change in both fuel types and in non-energy intermediates. Both fuel types and non-energy intermediates are inputs in total output. There are partial monopoly sellers of fuel-specific technologies, who hire specialized labor (“scientists”) to conduct research. The incentive for technical change is a market size effect – a larger market for a given fuel type provides higher returns to innovation and leads to more research. Discoveries are private for one period (five years, as in my model) and are subsequently available free of charge to other researchers.

The key finding of Fried (2018) is that in her calibrated model, induced technical change amplifies the effectiveness of a carbon tax by almost 20 percent. Specifically, she finds that when technological change is exogenous, a constant \$30 per ton CO2 carbon tax can reduce emissions by 30 percent in 20 years. But when technical change is endogenous the same outcome can be achieved with a tax of only \$24.5.

To test if my setting can deliver similar results, I modify my model to have a single region and two fuels – clean and dirty – analogous to Fried (2018). To replicate her calibration, I match parameters directly to their analogues in the data wherever possible. For example, the elasticities of substitution and demand have direct counterparts. Some parameters do not have a direct analogue. For example, the specification for the evolution of technology in Fried (2018) combines both investment in research and diffusion of past knowledge, producing returns to scale in research (and intra-fuel research spillovers) which are a function of the state. The equivalent objects in my model are constant, so I approximate by setting these to the value at the average growth rate. A third group of parameters are not explicitly reported. So I match these to IEA data for the US on fuel-specific consumption, prices, and emissions, on the grounds that both models should aim to reflect reality.³⁴ In line with Fried (2018), the green fuel is an aggregate of all zero-carbon energy sources, including nuclear energy. Table 3 describes the key parameter values and their rationale.

Parameter	Equivalent in Fried (2018)		
<i>Aggregate parameters</i>			
	<i>Value</i>		
Intra-fuel elasticity of substitution, \mathcal{E}	1.5		Identical
Elasticity of energy demand, F	0.05		Energy elasticity of substitution
Research function parameter, η	52.4		Matches returns to scale at avg. growth
Cross-fuel research spillovers, ρ_{12}, ρ_{21}	0.025		Matches spillovers at avg. growth
Research dampening parameter ω	1		Identical
<i>Fuel-specific parameters</i>			
	<i>Fossil</i>	<i>Green</i>	
Production function weights, λ_i	0.29	0.71	Match US energy production shares
Initial reference prices, $\bar{p}_{i,0}$	287	1000	Approx fuel prices, USD/ton oil equiv.
Carbon intensities, g_i	3.4	0	Match current US emissions

Table 3: Key model parameters to replicate Fried (2018)

I then subject this version of model to the same carbon tax policies as in Fried (2018). Table 4 compares the results. The headline responses of of emissions and energy usage are near-identical. This

³⁴Dollar fuel production costs matter for the conversion of the units of the carbon tax. The relative impact of a given tax per ton of CO2 is a decreasing function of the cost of production. When carbon taxes are zero, or for fuels which do not pollute, the level of input prices is irrelevant. In particular, the price of the green fuel has no bearing at all on the results.

is true both for the version where the endogenous growth channel is shut down and for the model with endogenous innovation.

The differences in some of the details below the top-level aggregate are a bigger, although some of this is a reflection of by the particular parameterization. For example, the response of the energy price is typically around half that in Fried. This is largely a result of the inelasticity of energy demand in the Fried model ($F = 0.05$). Given that the energy market adjusts principally on price rather than quantity, any differences in equilibrium in the energy market will show up only as differences in input prices. The low elasticity also explains why emissions fall so much despite energy usage remaining almost unchanged – the effect comes almost entirely through the composition channel. Similarly, differences in the relative usage and relative price of fuels likely stem from the fact that the calibration of the fuel-specific parameters does not map across identically between models. Small differences in initial fuel shares can translate into large differences in *percentage* changes in the impulse responses. That said, these differences also tend to dissipate over time and are typically much smaller in the long-run.

	Exogenous innovation (in 20 yrs.)		Endogenous innovation (in 20 yrs.)		Endogenous innovation (long run)	
	Model	Fried 2018	Model	Fried 2018	Model	Fried 2018
<i>Aggregates</i>						
Emissions	-31.8	-30.0	-29.8	-30.0	-37.3	-36.9
Energy usage	-0.5	-0.7	-0.4	-0.6	-0.4	-0.7
<i>Prices</i>						
Energy price	9.5	14.9	7.7	13.0	7.3	14.6
Green energy price	-0.0	0.6	-1.1	-1.0	-3.6	-2.6
<i>Green versus fossil</i>						
Relative price	-0.0	0.2	-3.7	-7.0	-12.1	-17.0
Relative usage	45.5	78.0	42.8	79.2	53.5	112.9
<i>Research input</i>						
Fossil technology	0.0	0.0	-27.2	-60.5	-28.8	-29.9
Green technology	0.0	0.0	11.9	53.3	12.6	23.8

Table 4: Percentage deviation from baseline in response to a permanent carbon tax, simplified model versus Fried (2018). Carbon taxes for exogenous and endogenous models are respectively \$30 and \$24.5 per ton CO₂

6.1.2 Comparison to Acemoglu et al. (2016)

Acemoglu et al. (2016) is a continuous time general equilibrium model, where households consume a continuum of products. Each product is produced by a partial monopolist who can choose either a clean or a dirty production technology. The productivities of the clean and dirty technologies are specific to each product, motivated by the idea that the way clean energy is utilized differs across products – transport services (electric cars and the like) potentially use renewable energy quite differently from large-scale steel production. Innovation is product specific, and follows a quality ladder, with the gains

of past research spilling over to future productivity. Emissions affect the stock of atmospheric CO₂, which in turn determines global temperatures. Higher temperatures lower aggregate productivity – a global climate externality.

The authors use this model to calculate optimal policies, balancing the distorting costs of two climate change mitigation policies – a carbon tax and a research subsidy – against the benefits of lower future temperatures. The key result in Acemoglu et al. (2016) is that the optimal policy in this setting is to aggressively subsidize research in the short term and increase carbon taxes later, shown in Figure A.5 in Appendix B.2.³⁵

I re-calibrate my model following a similar strategy to the preceding section. Parameters which map across directly are straightforward, including cross-fuel technology spillovers. Others require some interpretation. The research parameter η , again matches returns to scale in research at the average growth rate. And the elasticity of demand F is set to unity. As there is no separate energy sector in this model there is no direct analogue of F . However, the production function for each product is Cobb-Douglas, and so a uniform decrease productivity of both types of energy would cause a one percent increase in overall energy usage relative to the non-energy factor of production. This is mimicked in my model with $F = 1$.

The intra-fuel elasticity of substitution is even harder to compare across models. In Acemoglu et al. (2016), the substitution between clean and dirty modes of production happens at the individual product level. There, the elasticity is infinite as the producer just uses whichever technology is cheaper. The aggregate elasticity of substitution between clean and dirty fuel types is thus a function of the fraction of product manufacturers who are indifferent between the two energy types. As there is no way to directly capture a state-dependent elasticity of substitution such as this in the partial equilibrium model, I instead use the aggregate estimate of Papageorgiou et al. (2017), 3, on the grounds that Acemoglu et al. (2016) should be consistent with aggregate empirical facts. The remaining share parameters are again drawn from IEA data, this time for the whole world as Acemoglu et al. (2016) are a global model.

Acemoglu et al. (2016) use temperature as their main summary variable for the impact of the climate mitigation policies, and so I use this as the main point of comparison. To compute changes in temperature, I mirror exactly the mapping of Acemoglu et al. (2016) from emissions to temperature, who base theirs on Golosov et al. (2014).³⁶

Figure 11a shows the global temperatures under laissez-faire policies (taxes and subsidies of zero), in which case temperatures rise by almost 10 degrees Celsius over 200 years – a climate catastrophe by any standards. I then feed in to my model the optimal policies for the carbon tax and the research subsidy from Acemoglu et al. (2016). The resulting outcomes are almost identical, as shown in Figure

³⁵Acemoglu et al. (2016) use a carbon tax value tax, expressed in percent, which I convert to the more commonly used volume tax at equilibrium prices.

³⁶in this setting, temperature is proportionate to the log increase in the stock of atmospheric CO₂, with the coefficient of proportionality known as the climate sensitivity. Acemoglu et al. (2016) do not note this explicitly, but their results are consistent with the consensus median value of 3 (see, for example Meinshausen et al. (2009), Heal (2017), and Hassler et al. (2018)) which I also match. The stock of atmospheric CO₂ is a weighted average of past emission flows, with the weights depending on the speed at which atmospheric CO₂ decays or is absorbed by oceans and vegetation. For more details, readers are referred to Golosov et al. (2014).

11b. Temperatures across the two models vary by barely more than one tenth of one degree Celsius across the next 200 years, peaking at a little under 2 degrees Celsius above pre-industrial levels mid-way through the next century.

Whereas Fried (2018) shared many of the analytical building blocks with the partial equilibrium model, Acemoglu et al. (2016) is much less similar and so represents a much stiffer test. Nevertheless, the main aggregate responses to policy changes are much the same.

Parameter	Equivalent in Acemoglu et al. (2016)		
<i>Aggregate parameters</i>		<i>Value</i>	
Intra-fuel elasticity of substitution, \mathcal{E}	3	Central estimate in Papageorgiou et al. (2017)	
Elasticity of energy demand, F	1	Products Cobb-Douglas in inputs	
Research function parameter, η	33.0	Matches returns to scale at avg. growth	
Cross-fuel research spillovers, ρ_{12}, ρ_{21}	0	Identical	
Research dampening parameter ω	1	Identical	
<i>Fuel-specific parameters</i>		<i>Fossil</i>	<i>Green</i>
Production function weights, λ_i	0.39	0.61	Match global energy production shares
Initial reference prices, $\bar{p}_{i,0}$	433	1392	Global fuel prices, USD/ton oil equiv.
Carbon intensities, g_i	2.7	0	Match current global emissions

Table 5: Key model parameters equivalent to Acemoglu et al. (2016)




Figure 11: Temperature relative to pre-industrial average, Acemoglu et al. (2016) calibration.

In summary, the point of the two model-comparison exercises is to show that strong simplifying assumptions about the nature of endogenous technical change in Section 2 do not produce predictions for policy that are too far away from other approaches in the literature. It is therefore important to emphasize that this is an exercise in model testing, not model fitting. Throughout, I try to match the deep economic parameters of the micro-founded models. Where there is no direct mapping of the parameters across models, I match reasonable interpretations of them. And where that is impossible, or the requisite information is not reported, I match to the data. When I do this, the policy responses are very close to equivalent versions of my model. That this works for two very different models

studying different policies suggests that my approach, while not as rich as other model, is capturing the important major economic mechanisms in a quantitatively reasonable way.

6.2 Non-constant savings rate

A second important simplification in my model is that the savings rate is constant. This follows from the assumptions that (among other things) households have log preferences and that capital fully depreciates. These are quite strong assumptions, so one might reasonably wonder if the main policy implications remain if they are relaxed.

To investigate this concern, I re-solve the model with partial depreciation and CRRA preferences. This is computationally much more intensive than the solution with constant savings rates. This precludes re-computing the policies required to hit exactly the emissions targets for Europe and China. Instead, I take the policies computed there as given and verify ex post errors on emissions targets are small.

Specifically, I replace the household's problem with the following:

$$\begin{aligned} & \max \sum_{t=0}^{\infty} \beta^t \frac{C_{m,t}^{1-\gamma}}{1-\gamma} \\ \text{s.t.} \quad & C_{m,t} + K_{m,t+1} - (1 - \delta_K)K_{m,t} = w_{m,t}L_{m,t} + i_{m,t}K_{m,t} + \Omega_{m,t} \end{aligned}$$

Where δ_K is the depreciation rate of capital and $I_{m,t} = K_{m,t+1} - (1 - \delta_K)K_{m,t}$ is investment.³⁷ The Euler equation then becomes

$$\left(\frac{C_{m,t+1}}{C_{m,t}} \right)^\gamma = \beta (i_{m,t+1} + 1 - \delta_K) \quad (31)$$

Equilibrium in the extended model replaces equation (12) with (31). The resource constraint, equation (24), is replaced by:

$$C_{m,t} + I_{m,t} = Y_{m,t}$$

The resulting equilibrium savings rate is no longer constant, but instead has to be solved for in each region and time period to satisfy the Euler equation. Full details of the model solution are given in Appendix D.3.

I compute equilibrium under two policy scenarios: laissez-faire, and the policy settings which hit regional emissions targets in exercise (2) above (i.e. using carbon taxes to fund research subsidies). I set $\delta_K = 0.5$. With a five-year period this is equivalent to a 12 percent annual depreciation rate, an approximate mid-point of the BEA's estimates of depreciation on private nonresidential equipment. For the curvature of the utility function, I consider $\gamma = 2$ and $\gamma = 0.5$, an a way to cover the most commonly-used values. These changes necessitate recomputing the initial stocks of capital and labor productivity $K_{m,0}, A_{m,0}$ to hit the same calibration targets as the original model. These are that in

³⁷A further extension would be to let households' impatience, measured by β , to vary across countries. This would allow the model to match variation in savings rates across countries and would impact energy usage by increasing the capital stock in future in countries with higher savings rates. To some extent, though, the cross-regional variation in TFP growth rates already produces an effect of this type, with convergence in TFP growth acting as a proxy for catch-up growth via higher savings.




Figure 12: Change in relative to laissez-faire, effect of capital depreciation and non-log preferences

each region, 1) capital-output ratio is consistent with its long-run level and 2) the initial level of output still matches the data.

Overall, these changes have a negligible effect on the main findings of this paper. The policies computed by inverting the model emissions targets for China and OECD Europe still very nearly meet the emissions targets under the more general assumptions for depreciation and preferences (see Figure A.6 in Appendix B.2). Global outcomes are similarly unaffected, as shown in Figure 12; the impact of technological diffusion from emissions-reducing regions on both temperature and household consumption is almost identical across the different versions of the model.

Although capital accumulation is an important channel for the propagation of economic shocks and policy changes, these results suggest that it is near-irrelevant for climate mitigation policies.³⁸ The reason for this result is that there is limited scope for interaction between capital and the composition of energy. As discussed previously, the level of global emissions is determined by two factors: the amount of energy consumed and the composition of the inputs used in its production. Because the elasticity of substitution is considerably larger than the elasticity of demand, the latter channel is much more powerful than the former.³⁹ However, changes in the capital stock interact only total energy usage, independent of its composition, see equation (1). Thus, capital has only a minor effect emissions because it affects only the more minor determinant of emissions – the quantity of energy produced – and has no bearing at all on the most quantitatively important determinant of emissions – the composition of fuel inputs.

Although the foregoing suggests that when the savings rate is endogenous, capital has very little effect on the environment, the reverse is not true. By eliminating the climate externality, climate mitigation policies raise the marginal productivity of capital in the long-run. And because the Euler equation fixes the returns of capital in the very long run (the oft-cited mantra that “the long-run

³⁸This result is not unique to this setting – Barrage (2013) finds that optimal global climate mitigation policies do not depend crucially on assumptions about capital accumulation.

³⁹See discussion in section 4.3.1.

elasticity of capital is infinite”) climate mitigation policies cause the savings rate to fall slightly in the long-run – by around one percentage point. This is not welfare-reducing, though. With higher total factor productivity, households can afford to consume more, both in the present and – because the saved capital is subsequently more productive – in future.

7 Conclusions

The main thrust of this paper is to investigate the extent to which regional pursuit of emissions targets might provide a solution to global climate change via the international diffusion of induced technological change. The central result is that in a model which reflects standard views in the literature on the path of future emissions, the speed, and key economic elasticities, that this effect might be large enough for successful emissions-reduction policies in China and OECD Europe to avoid the worst-case spiraling catastrophic climate change. This is robust to assumptions about the speed and patterns of technological diffusion, determinants of the savings rate, and the nature of endogenous technical change. Beyond this core result, I also provide explanations of how different policies work and the mechanism for international technological diffusion.

The central result is good news for the global environment in an obvious and direct sense. It means that failure to achieve full global cooperation on emissions limits need not condemn the planet to an inevitable climate catastrophe. The international diffusion of new technologies (and, just as importantly the non-diffusion of new dirty technologies) could provide a way to avoid the very worst scenarios even if we cannot achieve full global cooperation. Given that progress towards such global cooperation over the last three decades has proved at best elusive (and at worst, an abject failure) this assessment is certainly welcome.

There is also a more indirect sense in which the central result is good news. Positive international policy spillovers make individual climate mitigation programs more appealing because technological diffusion amplifies their effect on global emissions. As a result, the old argument for individual inaction – that limiting emissions is individually costly and does little to address a global problem – is weakened. Of course, other parts of the world are still free-riding on the economic sacrifices of countries which implement emissions reductions. But at least that sacrifice is to some end – avoiding the very worst forms of global climate change. And so the very fact that individual action might be more appealing than otherwise imagined could mean that action by major polluters becomes more likely. This also has a bearing on the stability of global agreements to reduce emissions. Because clean innovations in different regions partially crowd each other out, technological spillovers cause the *marginal* gain of staying in a climate agreement to increase as participation declines.⁴⁰ And so international agreements to reduce emissions are less likely to unravel completely if one or two major polluters fail to participate or do not follow through on their commitments.

A final important point for discussion is to ask what possible further issues not captured in this analysis might threaten or undermine the the central findings. Two obvious items spring to mind.

⁴⁰In the model, this is reflected in the fact that the fuel-specific minimum price determines global input costs. More intuitively, the marginal gain from a region implementing policies which result in, say, an improved solar panel is diluted if other regions also have similar policies which induce similar innovations – something can only be invented once.

The first is trade, which – other than oil – is entirely excluded from this analysis and might have an impact in a variety of ways. For instance, the opportunity to sell new technologies – both clean and dirty – overseas might be an important driver of domestic growth in emitters with large export shares (e.g. Germany or China). Moreover, overseas demand might be also an important determinant of research effort. If demand elsewhere, e.g. in the United States, drives research effort in Europe and China then policies which shift domestic markets towards green energy supplies might not induce as large a switch in research effort as in the model. However, this channel will also have a countervailing effect – research in places other than Europe and China would likely switch towards green technologies with a view to exporting the associated products to those regions. The overall impact of these effects is not obvious.

The second potentially important extension is fuel-specific capital. As noted in Section 6, the inclusion of long-lived capital does not substantially effect the main results of this paper. In the model, capital is combined with labor and then this bundle is combined with energy. In reality, though, energy production itself demands capital investments, often specific to the type of fuel used – coal-fired power plants, gas stations, and the like. Abandoning investments in such “stranded assets” might be rather costly, slowing the reallocation of energy inputs from high- to low-carbon fuels. One way to proxy for this rather simply might be consider the impact of lowering the intra-fuel elasticity of substitution \mathcal{E} . This would necessitate more aggressive policies to meet emissions targets (i.e. higher carbon taxes and subsidies), but conditional on meeting emissions reductions, the technological effect would be the same. This suggests that the main impact of fuel-specific capital would be to make meeting emissions targets more costly, but would be otherwise minimal. Further exploration of both this issue and impact of trade are left for future work.

References

- Acemoglu, D. (2002). Directed technical change. *The Review of Economic Studies* 69(4), 781–809.
- Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The environment and directed technical change. *American Economic Review* 102(1), 131–66.
- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016). Transition to clean technology. *Journal of Political Economy* 124(1), 52–104.
- Barrage, L. (2013). Sensitivity analysis for golosov, hassler, krusell, and tsyvinski (2013): Optimal taxes on fossil fuel in general equilibrium. *Technical Notes*.
- Barrage, L. (2020). Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy. *The Review of Economic Studies* 87(1), 1–39.
- Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The review of economic studies* 83(1), 87–117.
- Bosetti, V., C. Carraro, and M. Galeotti (2006). The dynamics of carbon and energy intensity in a model of endogenous technical change. *The Energy Journal* (Special Issue# 1).
- Buera, F. J. and E. Oberfield (2020). The global diffusion of ideas. *Econometrica* 88(1), 83–114.
- Buonanno, P., C. Carraro, and M. Galeotti (2003). Endogenous induced technical change and the costs of kyoto. *Resource and Energy economics* 25(1), 11–34.
- Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature* 527(7577), 235–239.
- Caselli, F. and W. J. Coleman (2001). Cross-country technology diffusion: The case of computers. *American Economic Review* 91(2), 328–335.
- Comin, D. and B. Hobijn (2010). An exploration of technology diffusion. *American economic review* 100(5), 2031–59.
- Council, D. P. (2013). Technical support document:-technical update of the social cost of carbon for regulatory impact analysis-under executive order 12866. *Environmental Protection Agency*.
- Dechezleprêtre, A., M. Glachant, and Y. Ménière (2013). What drives the international transfer of climate change mitigation technologies? empirical evidence from patent data. *Environmental and Resource Economics* 54(2), 161–178.
- Eaton, J. and S. Kortum (1999). International technology diffusion: Theory and measurement. *International Economic Review* 40(3), 537–570.
- Fried, S. (2018). Climate policy and innovation: A quantitative macroeconomic analysis. *American Economic Journal: Macroeconomics* 10(1), 90–118.
- Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82(1), 41–88.
- Hassler, J., P. Krusell, and C. Olovsson (2018). The consequences of uncertainty: climate sensitivity and economic sensitivity to the climate. *Annual Review of Economics* 10, 189–205.

- Hassler, J., P. Krusell, C. Olovsson, and M. Reiter (2020). On the effectiveness of climate policies. *Working paper, IIES Stockholms universitet*.
- Heal, G. (2017). The economics of the climate. *Journal of Economic Literature* 55(3), 1046–63.
- Meinshausen, M., N. Meinshausen, W. Hare, S. C. Raper, K. Frieler, R. Knutti, D. J. Frame, and M. R. Allen (2009). Greenhouse-gas emission targets for limiting global warming to 2 c. *Nature* 458(7242), 1158–1162.
- Nordhaus, W. (2010). Excel file for rice model as of april 26, 2010.
- Nordhaus, W. (2014a). Estimates of the social cost of carbon: concepts and results from the dice-2013r model and alternative approaches. *Journal of the Association of Environmental and Resource Economists* 1(1/2), 273–312.
- Nordhaus, W. (2014b). *A question of balance: Weighing the options on global warming policies*. Yale University Press.
- Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the” dice” model. *The American Economic Review* 83(2), 313–317.
- Nordhaus, W. D. (2014c). The perils of the learning model for modeling endogenous technological change. *The Energy Journal* 35(1).
- Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences* 114(7), 1518–1523.
- Olivier, J. G., K. Schure, and J. Peters (2017). Trends in global co2 and total greenhouse gas emissions. *PBL Netherlands Environmental Assessment Agency* 5.
- Pachauri, R. K., M. R. Allen, V. R. Barros, J. Broome, W. Cramer, R. Christ, J. A. Church, L. Clarke, Q. Dahe, P. Dasgupta, et al. (2014). *Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change*. Ipcc.
- Papageorgiou, C., M. Saam, and P. Schulte (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics* 99(2), 281–290.
- Popp, D., R. G. Newell, and A. B. Jaffe (2010). Energy, the environment, and technological change. In *Handbook of the Economics of Innovation*, Volume 2, pp. 873–937. Elsevier.
- Sampson, T. (2016). Dynamic selection: an idea flows theory of entry, trade, and growth. *The Quarterly Journal of Economics* 131(1), 315–380.
- Stern, N. and N. H. Stern (2007). *The economics of climate change: the Stern review*. cambridge University press.
- Tol, R. S. (2002). Estimates of the damage costs of climate change. part 1: Benchmark estimates. *Environmental and resource Economics* 21(1), 47–73.
- Van den Bergh, J. C. and W. Botzen (2015). Monetary valuation of the social cost of co2 emissions: a critical survey. *Ecological Economics* 114, 33–46.
- Van der Werf, E. (2008). Production functions for climate policy modeling: An empirical analysis. *Energy economics* 30(6), 2964–2979.

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

A The structure of the energy market

Here I show two results about the energy market: that the average cost function is downward-sloping; and that equilibrium exists if $F < \eta$ and the government caps prices. Throughout, I suppress dependence on the region m and set aside the special status of oil, assuming for simplicity that all inputs are improvable.

A.1 The average cost function

The average cost function is the per unit price of fuel inputs and research expenditure:

$$\begin{aligned} \frac{C(E_t)}{E_t} &= V_t + \frac{R_t}{E_t} \\ &= \frac{\sum_{i=1}^N \hat{p}_{i,t} e_{i,t}}{E_t} + \frac{\sum_{i=1}^N p_{i,t} e_{i,t}}{(\eta - 1)E_t} \end{aligned}$$

Where

$$V_t = \left(\sum_{i=1}^N \lambda_i^{\frac{1}{\rho-1}} \hat{p}_{i,t}^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}} = \frac{\sum_{i=1}^N \hat{p}_{i,t} e_{i,t}}{E_t}$$

is the CES price index for energy inputs.

$$C(E_t) = \sum_{i=1}^N \hat{p}_{i,t} e_{i,t} + \frac{1}{\eta - 1} \sum_{i=1}^N e_{i,t} p_{i,t}$$

When $\tau_t = 0$ the (absolute) slope of the cost function will be maximized. This is because as τ_t increases, the share of pre-tax costs in the total decline. As research is proportional to pre-tax costs then research too becomes a smaller share of the total cost. Because the fixed research cost is the source of the slope in the aggregate cost function $C(E_t)$, then as τ_t grows ever large, the cost function gets ever straighter, limiting to $C(E_t) = V_t E_t$. The slope when $\tau_t = 0$ therefore puts an upper bound on the magnitude of the slope of the cost function, and so I focus on this case.

When $\tau_t = 0$, $\hat{p}_{i,t} = p_{i,t}$, and we get that:

$$C(E_t) = \frac{\eta}{\eta - 1} \sum_{i=1}^N e_{i,t} p_{i,t}$$

Then using $d \log Y_t$ as shorthand for $d \log Y_t / d \log E_t$, we get that:

$$\begin{aligned} d \log C(E_t) &= \sum_{i=1}^N w_{i,t} d \log p_{i,t} + \sum_{i=1}^N w_{i,t} d \log e_{i,t} \\ \Rightarrow d \log (C(E_t)/E_t) &= \sum_{i=1}^N w_{i,t} d \log p_{i,t} \end{aligned}$$

Where $w_{i,t} = p_{i,t} e_{i,t} / \sum_{i=1}^N p_{i,t} e_{i,t}$ is the expenditure share of input i . The second line follows from

the properties of CES production, that $d \log E_t = \sum_{i=1}^N w_{i,t} d \log e_{i,t}$.

To figure out $d \log p_{i,t}$, I follow a guess-and-verify strategy, observing that the production function is homothetic. I guess that relative prices all remain unchanged. Then hometheticity means that all inputs all increase in line with E_t , i.e. $d \log e_{i,t} = 1 \quad \forall i$.

Equation (9) can be differentiated to give:

$$d \log e_{i,t} = -\eta d \log p_{i,t} \quad \forall i$$

Thus, if $d \log e_{i,t} = 1$ were true, then $d \log p_{i,t} = -1/\eta \quad \forall i$. This verifies the guess that relative prices remain the same. Research is also homothetic here; when total demand E goes up by 1%, inputs all increase by 1% and prices fall by $1/\eta\%$ and increasing the overall cost by $1 - 1/\eta\%$.

Thus,

$$d \log (C(E_t)/E_t) = -\frac{1}{\eta}$$

Or, using that $d \log E_t = 1$:

$$d \log (C(E_t)) = 1 - \frac{1}{\eta}$$

A.2 Market structure and pricing

With the minimized unit cost function $C(\cdot)$ in hand, we can return to profit maximization. This is because profit maximization can be broken into two stages: cost-minimization conditional on output E_t – which determined the composition of energy inputs – and maximization of profit over E_t conditional on cost minimization – which determines the scale of production.

$$\max_{E_t} \Pi_t = P_t(E_t)E_t - C(E_t) \quad \text{s.t. } \Pi \geq 0 \quad (32)$$

I focus first on the case where $\tau_t = 0$ and derive the conditions under which equilibrium exist. Then it is relatively straightforward to argue that these conditions guarantee equilibrium for $\tau_t > 0$.

Integrating the derivatives of the cost function above, we get that

$$C(E) = A^S E^{1-\frac{1}{\eta}} \quad (33)$$

Where A^S is a constant. Under monopoly, the firm's demand curve is just the aggregate energy demand curve. From the firm's perspective this is just a CES demand curve with elasticity F ,

$$E_t = A_t^D (P_t)^{-F} \quad (34)$$

where A_t^D is a demand shifter, determined in general equilibrium. Then the monopolist's profit function is

$$\Pi_t(E) = A_t^D E^{1-\frac{1}{F}} - A^S E^{1-\frac{1}{\eta}}$$

Which we can differentiate to get:

$$\Pi'_t(E_t) = \left(1 - \frac{1}{F}\right) A_t E_t^{-\frac{1}{F}} - \left(1 - \frac{1}{\eta}\right) A_t^S E_t^{-\frac{1}{\eta}}$$

It is clear that if $F > \eta$ then profit increases forever as $E_t \rightarrow \infty$, so we must have that $F < \eta$. If $1 < F < \eta$ then they have a concave profit function and a well-defined interior solution. But the standard parameterization in the literature is for energy demand to be inelastic, so $F < 1$. So if $F < \eta$ and $F < 1$ then the derivative of the profit function is always negative and profits are maximized in the limit as $E_t \rightarrow 0$. Thus, we must have some sort of regulation to cap prices and guarantee an interior solution.

For $\tau_t > 0$, the same argument goes through with whatever the elasticity of the cost function is. Denote this elasticity by $1/\eta'$. Because we know that the (absolute) elasticity of the cost function is maximized at $\tau_t = 0$, then we must have that $\eta' > \eta$. And thus if $F < \eta$, then $F < \eta'$ as well. So the same conditions for equilibrium hold here.

B Further details of the model

B.1 Calibration details

	α_m	$\log \nu_m$	$\log K_{m,0}$	$\log A_{m,0}$
Asia Pacific	0.41	-18.66	14.18	10.13
China	0.41	-18.68	15.27	10.52
European OECD	0.39	-20.62	15.61	11.95
India	0.41	-17.48	13.66	9.68
Japan	0.49	-20.98	14.49	12.03
Latin America and Carribean	0.41	-19.09	14.02	10.66
Rest of OECD	0.51	-19.55	14.90	11.70
Rest of World	0.41	-17.23	13.11	9.97
Russia	0.41	-16.83	13.15	11.10
Sub-Saharan Africa	0.41	-17.22	12.78	9.49
USA	0.38	-20.39	15.63	12.41

Table A.1: Region-specific calibration parameters and initial conditions.

	Conv. Oil	Coal	Dom. Oil	Gas	Hydro	Nuclear	Renewables
Asia Pacific	0.33	0.09	0.24	0.23	0.23	0.17	0.88
China	0.41	0.28	0.29	0.20	0.58	0.36	0.75
European OECD	0.34	0.10	0.21	0.24	0.28	0.55	0.60
India	0.35	0.14	0.15	0.12	0.24	0.19	0.96
Japan	0.60	0.12	0.02	0.42	0.31	0.19	0.53
Latin America and Carribean	0.03	0.04	0.35	0.20	0.49	0.15	0.79
Rest of OECD	0.26	0.12	0.36	0.22	0.47	0.50	0.57
Rest of World	0.03	0.09	0.39	0.35	0.33	0.39	0.69
Russia	0.02	0.10	0.41	0.46	0.36	0.57	0.26
Sub-Saharan Africa	0.14	0.07	0.12	0.05	0.17	0.11	0.93
USA	0.37	0.13	0.34	0.15	0.25	0.69	0.63

Table A.2: Fuel-specific productivities, $\lambda_{m,i}$.

	Conv. Oil	Coal	Dom. Oil	Gas	Hydro	Nuclear	Renewables
Asia Pacific	2.32	3.99	2.32	2.12	0.00	0.00	0.00
China	2.37	3.79	2.37	2.21	0.00	0.00	0.00
European OECD	2.54	3.92	2.54	2.26	0.00	0.00	0.00
India	2.53	3.80	2.53	1.37	0.00	0.00	0.00
Japan	2.40	3.87	2.40	2.39	0.00	0.00	0.00
Latin America and Carribean	2.76	3.92	2.76	2.04	0.00	0.00	0.00
Rest of OECD	2.42	3.95	2.42	2.26	0.00	0.00	0.00
Rest of World	2.85	3.88	2.85	2.17	0.00	0.00	0.00
Russia	1.94	3.51	1.94	2.09	0.00	0.00	0.00
Sub-Saharan Africa	3.04	3.60	3.04	1.70	0.00	0.00	0.00
USA	2.52	3.94	2.52	2.25	0.00	0.00	0.00

Table A.3: Pollution intensities, tonnes CO2 per tonne oil equivalent, $g_{m,i}$.

	Coal	Dom. Oil	Gas	Hydro	Nuclear	Renewables
Asia Pacific	143.78	465.47	371.17	1360.00	1082.98	1469.39
China	156.85	475.82	370.05	1468.06	1152.29	1466.01
European OECD	146.24	508.18	383.94	1427.59	1216.51	1476.35
India	147.89	447.77	353.26	1360.54	1086.90	1472.99
Japan	121.59	448.61	489.45	1336.43	1048.34	1358.93
Latin America and Carribean	134.36	475.90	365.40	1432.62	1067.85	1452.12
Rest of OECD	171.92	420.09	280.19	1431.29	1167.38	1422.28
Rest of World	140.62	470.46	373.56	1366.97	1122.36	1410.94
Russia	141.69	472.36	380.87	1373.04	1153.43	1314.83
Sub-Saharan Africa	141.03	441.76	333.86	1329.98	1048.02	1472.51
USA	175.97	408.07	154.59	1392.38	1214.68	1454.99

Table A.4: Reference input costs, 2012 USD per tonne oil equivalent, $\bar{g}_{m,i}$.

B.2 Extra charts




Figure A.1: Regional outcomes under laissez-faire.

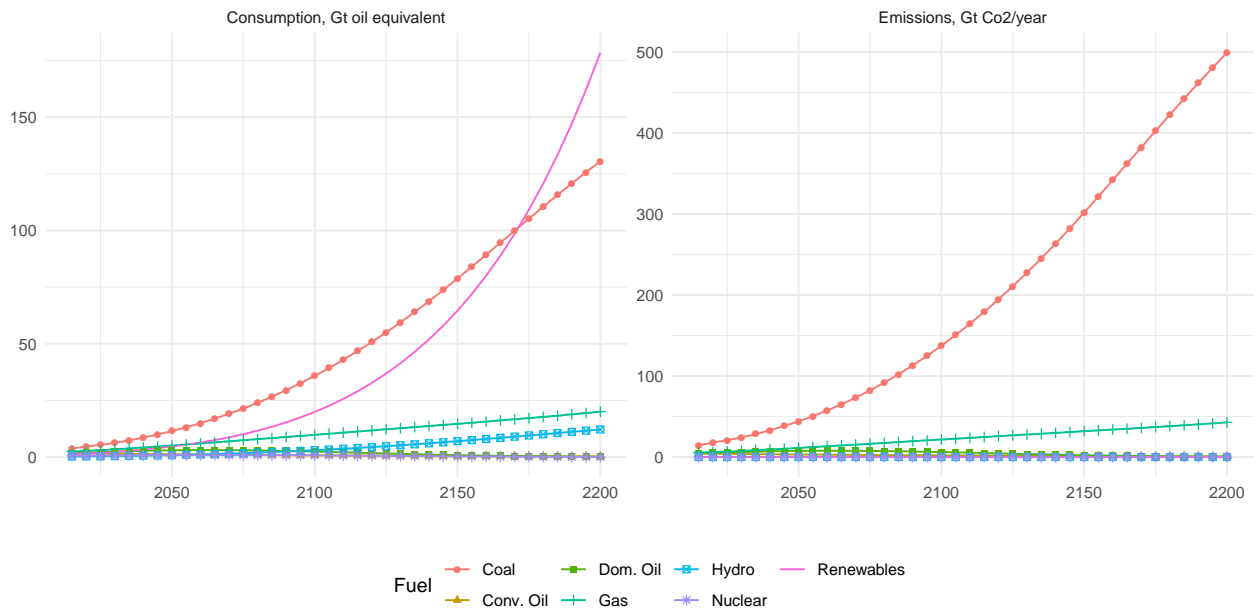


Figure A.2: Fuel consumption and emissions under laissez-faire.

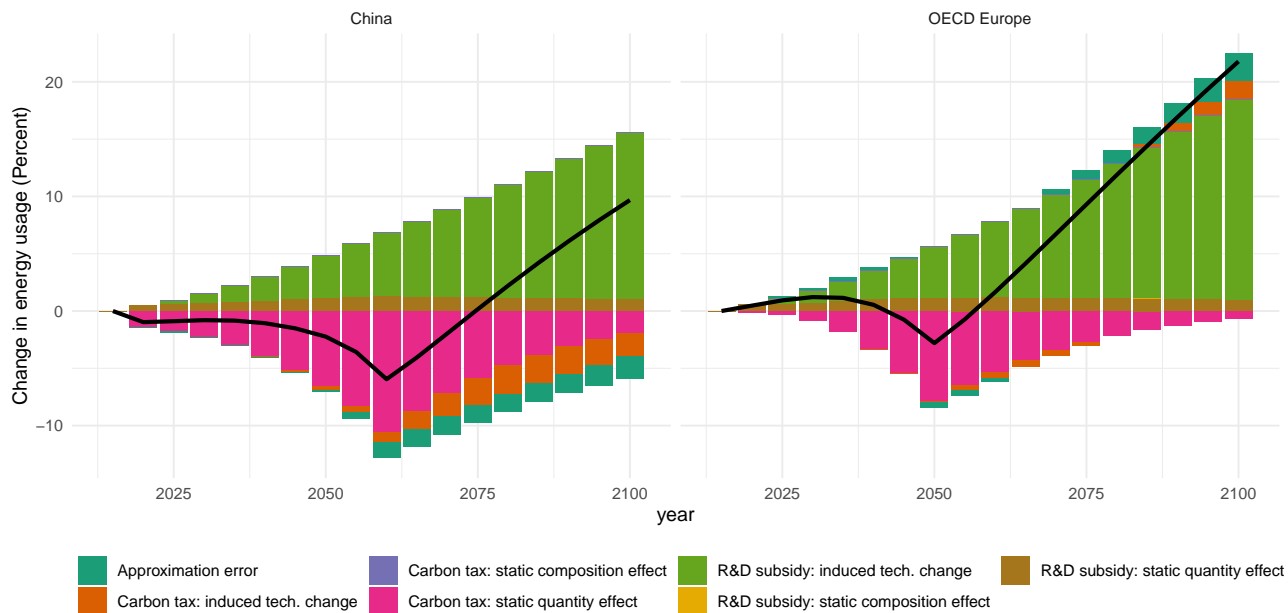


Figure A.3: Policy decomposition of changes in emissions relative to laissez-faire. Using analytical impulse responses from a partial equilibrium approximation to the model, the impact of each policy is divided into three channels: static composition and quantity effects, and the dynamic induced technical change effect.




Figure A.4: Household consumption, relative to laissez-faire. Numbers in parenthesis refer to policy experiments, consistent across charts.






Figure A.5: Optimal policies in Acemoglu et al. (2016).



(a) China



(b) OECD Europe

Figure A.6: Emissions versus target, effect of alternate assumptions for the determinants of the savings rate

C A partial equilibrium approximation

Here I outline a simplified version of the model which studies equilibrium in only the energy market. The aim here is to understand why particular policies can achieve emissions targets. This can be somewhat difficult in the main model as there are so many moving parts. So here I strip away enough of these that we can compute exactly the impact of policy on emissions and energy usage. One can also decompose the analytical impulse responses into three channels – the static composition and quantity effects, and dynamic induced technical change. Of course, this comes at a cost. These simplification will mean the impulse responses are not numerically identical to those in the full general equilibrium model. But as Figures 7 and A.3 show, this cost is small, as the approximation errors are negligible.

Specifically, this setting studies only one market – energy – in one region, suppressing general equilibrium effects from energy demand, climate externalities, and international technology spillovers. Given the importance of spillovers on regions which do not reduce emissions, this is not a useful vehicle for analyzing what happens outside China and OECD Europe. But as the impact on emissions-reducing regions comes principally from their own policies, this is a useful simplification for understanding what is happening there.

In the following, I first outline the dynamic system of equations defining the partial equilibrium model, then calculate the impulse responses by differentiating this system of equations with respect to the policies: $\tau_t, \delta_{j,t}, \chi_{j,t}$. This proceeds in three stages. First, tracing through the impact on the period t allocation. Second, calculating how reference prices in period $t+1, t+2, \dots$ respond to policy changes in period t . Third, describing how changes in reference prices affect allocations in those periods.

C.1 Model equations

Here, I present the equations defining a partial equilibrium approximation to the main model. Because this is a single-region model, I suppress dependence on m throughout.

Firm's first order conditions for cost-minimization are identical to those in the main text, but with one exception. I assume that oil is the same as all other fuel types and can be improved. This difference will be subsumed in the approximation error. For the sake of completeness, these are repeated here:

$$\frac{e_{i,t}}{E_t} = \left(\frac{\hat{p}_{i,t}}{V_t \lambda_i} \right)^{-\epsilon} \quad (35)$$

$$(1 - \delta_{i,t}) e_{i,t} = \frac{\epsilon_i (1 - \chi_{i,t})}{\bar{p}_{i,t}} \left(\frac{p_{i,t}}{\bar{p}_i} \right)^{-\eta} \quad (36)$$

$$r_{i,t} = \left(\frac{1 - \delta_{i,t}}{\eta - 1} p_{i,t} \right) e_{i,t} \quad (37)$$

Where:

$$\hat{p}_{i,m,t} = \tau_{m,t} g_i + (1 - \delta_{i,m,t}) p_{i,t} \quad (38)$$

I assume that energy demand has constant elasticity, F . That is:

$$E_t = A_t (P_t)^{-F} \quad (39)$$

Where A_t is an exogenous demand shifter. The exogeneity of A_t is one of the ways that this differs from a general equilibrium setting. In general equilibrium, shifts in the demand curve are potentially endogenous to outcomes in the energy market. For example, the price of energy can affect real household incomes, impacting their demand for all goods (including energy).

Appendix A outlines the conditions under which the preceding cost function and demand curve give rise to a natural monopoly. I assume that these conditions hold and that the energy market is regulated so that firms make zero profit. Thus, price equals average cost:

$$P_t = V_t + \frac{R_t}{E_t} \quad (40)$$

Where, as before, V_t is the CES input price aggregator

$$V_t = \left(\sum_{i=1}^N \lambda_i^{\frac{1}{\rho-1}} \hat{p}_{i,t}^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}} = \frac{\sum_{i=1}^N \hat{p}_{i,t} e_{i,t}}{E_t}$$

Technological progress diffuses over time and across fuels

$$\log \bar{\mathbf{p}}_{t+1} = (I_N - \omega\rho) \log \bar{\mathbf{p}}_t + \omega\rho \log \mathbf{p}_t \quad (41)$$

Where $\bar{\mathbf{p}}_t$ and \mathbf{p}_t are the vectors of reference and realized production costs, ρ is the matrix with row i equal to ρ_i and I_N is the $N \times N$ identity matrix.

Equilibrium in this simplified model is given by a sequence of fuel specific prices, quantities and research expenditures $\{p_{i,t}, \bar{p}_{i,t}, \hat{p}_{i,t}, e_{i,t}, r_{i,t}\}$ and aggregate prices and quantities P_t, E_t which, given initial reference prices $\bar{p}_{i,t}$ and sequence of exogenous demand shifters A_t , satisfy equations (35)-(41).

C.2 Static derivatives

For all policies, the responses of the demand curve and pricing equations are the same. To simplify notation, I suppress dependence on t and use dY to simply mean the response of variable Y to any contemporaneous policy change.

The demand curve is:

$$d \log E = -F d \log P \quad (42)$$

From equation (40), we have that:

$$d \log P = w_P d \log V + (1 - w_P) d \log R - (1 - w_P) d \log E \quad (43)$$

Where $w_P = V/P$ is the input share of the final price. Using the first order conditions, we have

that:

$$d \log V = \sum_{i=1}^N w_i d \log \hat{p}_i \quad (44)$$

$$d \log R = \sum_{i=1}^N v_i (d \log e_i + d \log p_i) \quad (45)$$

Substituting equations (43), (44), and (45) into (42) we get that:

$$\sum_{i=1}^N w_i d \log e_i = -F w_P \sum_{i=1}^N w_i d \log \hat{p}_i - F(1 - w_P) \sum_{i=1}^N v_i d \log p_i + F(1 - w_P) \sum_{i=1}^N (w_i - v_i) d \log e_i \quad (46)$$

This is a key equation in all the policy derivatives, as it imposes the constraint that changes in the sum of prices and quantities must be consistent with the aggregate demand curve

C.2.1 Carbon tax

In this subsection d represents the derivative w.r.t. a one dollar increase in the carbon tax, i.e. $d/d\tau$. As before, I will also assume that good j is the clean energy fuel and use notation \tilde{x}_i to represent the log difference relative to clean energy, i.e. $\tilde{x}_i = \log x_i - \log x_j$.

Changes in energy composition. Differentiating the first order conditions for factor inputs we get that:

$$d\tilde{e}_i = -\mathcal{E} d\tilde{p}_i \quad (47)$$

$$d \log e_i = -\eta d \log p_i \quad \forall i \quad (48)$$

We can rewrite (48) as:

$$d\tilde{e}_i = -\eta d\tilde{p}_i \quad \forall i \neq j \quad (49)$$

In the carbon tax case, the post-tax price differential can be rewritten as:

$$d\tilde{p}_i = d\tilde{p}_i + \frac{g_i}{\hat{p}_i} - \frac{g_i \tau}{\hat{p}_i} d \log p_i \quad (50)$$

Solving for the system of equations (47), (49), (50), we get that:

$$d\tilde{p}_i = \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right) \left(\frac{g_i}{\hat{p}_i} \right) (1 - \tau d \log p_i)$$

Which we can rewrite as:

$$d\tilde{p}_i = H_i (1 - \tau d \log p_i) \quad (51)$$

Where $H_i = \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right) \left(\frac{g_i}{\hat{p}_i} \right)$. We can further use the definition of $d\tilde{p}_i$ to replace $d \log p_i$ with the clean energy equivalent.

$$d\tilde{p}_i = Q_i(\tau) (1 - \tau d \log p_j) \quad (52)$$

Where $Q_i(\tau) = (1/H_i + \tau)^{-1}$. Note that we can use (47) and (49) to also solve easily for the quantities

and post-tax price equivalents:

$$d\tilde{e}_i = -\eta Q_i(\tau)(1 - \tau d \log p_j) \quad (53)$$

$$d\tilde{p}_i = \frac{\eta}{\mathcal{E}} Q_i(\tau)(1 - \tau d \log p_j) \quad (54)$$

Note that compositional changes are therefore initially proportionate to the social cost of pollution, $\frac{g_i}{p_i}$, but that this is mitigated (for given p_j) as τ increases, simply because the tax itself becomes an increasingly large share of the post-tax price. The role of technical change is as an amplifying mechanism, here captured by the term $\mathcal{E}/(\eta - \mathcal{E})$.

Changes in energy usage and emissions. We now rewrite (46) in terms of changes relative to the clean fuel, good j :

$$\begin{aligned} d \log e_j + \sum_{i=1}^N w_i d\tilde{e}_i &= -F(w_P d \log \hat{p}_j + (1 - w_P) d \log p_j) - F w_P \sum_{i=1}^N w_i d\tilde{p}_i \\ &\quad - F(1 - w_P) \sum_{i=1}^N v_i d\tilde{p}_i + F(1 - w_P) \sum_{i=1}^N (w_i - v_i) d\tilde{e}_i \end{aligned} \quad (55)$$

Collecting like terms, we have that:

$$d \log p_j = \frac{Z(w_P, \tau)}{F - \eta + \tau Z(w_P, \tau)} \quad (56)$$

Where:

$$Z(w_P, \tau) = \left(\eta - F + F w_P \left(\frac{\mathcal{E} - \eta}{\mathcal{E}} \right) \right) \bar{Q}(\tau) - \eta F(1 - w_P)(\bar{Q}(\tau) - \bar{Q}^0(\tau))$$

And $\bar{Q}(\tau) = \sum_{i=1}^N w_i Q_i(\tau)$ and $\bar{Q}^0(\tau) = \sum_{i=1}^N v_i Q_i(\tau)$ are the averages of Q_i using the pre- and post-tax input weights. Note of course that w_P itself is an equilibrium variable which will depend on τ , so this is really just a function in τ , albeit indirectly. From this, we can solve immediately for e_i, p_i from equations (53) and (52) and the identity $d\tilde{x}_i = d\tilde{x}_i + dx_j$.

Note that as τ increases, $d \log p_j$ tends to $1/\tau$. In other words, because tax starts to be an ever-increasing fraction of the price that despite being a volume tax, the carbon tax starts to look ever more like a value tax. The discrimination between different polluting fuel types vanishes, suppressing the strong compositional changes induced by the first dollar of the tax.

To compute the marginal response for aggregate variables at any τ , we now need only follow the following four substitutions:

1. Using definitions of $d\tilde{x}_i$ and rearranging equation (52) we get that:

$$d \log p_i = Q_i(\tau) + (1 - \tau Q_i(\tau)) d \log p_j$$

Substituting in for $d \log p_j$ using equation (56) then gives $d \log p_i$

2. From this

3. Derivatives for E_t , M_t then follow from:

$$d \log M_t = \sum_{i=1}^N w_{i,t}^M d \log e_{i,t}$$

$$d \log E_t = \sum_{i=1}^N w_{i,t} d \log e_{i,t}$$

Where $w_{i,t}, w_{i,t}^M$ are the shares of fuel i in expenditures and emissions respectively (so $w_{i,t}^M = g_i e_{i,t} / M_t$).

	Coal	Nat. Gas	Oil	Zero-Carbon	Total
<i>Data</i>					
Energy, $e_{i,t}$ (MtOE)	1956	165	539	315	2976
Price, $p_{i,t}$ (\$ per TOE)	149	367	513	1364	
Emissions, $m_{i,t}$ (Mt CO2)	7409	364	1280	0	9053
<i>Shares and ratios</i>					
Emissions intensity, g_i (tonnes CO2 per TOE)	3.79	2.21	2.37	0.00	
Expenditure share, $w_{i,t}$	0.28	0.06	0.26	0.41	
Emissions share, $w_{i,t}^m$	0.82	0.04	0.14	0.00	
Social cost, $g_{i,t} / \hat{p}_{i,t}$ (tonnes CO2 per \$)	0.025	0.006	0.005	0.000	
Avg. social cost, $\sum_i w_{i,t} g_i / \hat{p}_{i,t}$	0.009	0.009	0.009	0.009	
<i>Marginal responses</i>					
Energy consumption, $de_{i,t} / d\tau_t$, (pp, contrib.)	-1.1	0.0	0.1	0.6	-0.3
Composition effect	-1.0	0.0	0.2	0.7	-0.0
Quantity effect	-0.1	-0.0	-0.1	-0.1	-0.3
Change in emissions, $dm_{i,t} / d\tau_t$, (pp, contrib.)	-3.2	0.0	0.1	0.0	-3.1
Composition effect	-3.0	0.0	0.1	0.0	-2.8
Quantity effect	-0.2	-0.0	-0.0	-0.0	-0.3
<i>Memo: Research amplification factors</i>					
Energy composition, $\eta / (\eta - \mathcal{E})$					1.07
Total energy consumption, $\eta / (\eta - F)$					1.01

Table A.5: Static marginal impact of the carbon tax at $\tau_t = 0$, worked example using averages 2015-2019. Marginal responses expressed as percentage point contributions to total, China.

In the case where $\tau = 0$, things simplify even further to give:

$$\frac{d \log p_{i,t}}{d\tau_{i,t}} = \mathcal{E} \left(\frac{1}{\eta - \mathcal{E}} \right) \left(\frac{g_i}{\hat{p}_{i,t}} - \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}} \right) + F \left(\frac{\eta - 1}{\eta} \right) \left(\frac{1}{\eta - F} \right) \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}}$$

$$\frac{d \log e_{i,t}}{d\tau_{i,t}} = \underbrace{-\mathcal{E} \left(\frac{\eta}{\eta - \mathcal{E}} \right) \left(\frac{g_i}{\hat{p}_{i,t}} - \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}} \right)}_{\text{Change in energy composition}} - \underbrace{F \left(\frac{\eta - 1}{\eta} \right) \left(\frac{\eta}{\eta - F} \right) \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}}}_{\text{Change in total energy consumption}}$$

Which can be used to derive the responses of energy and emissions:

$$\begin{aligned}
\frac{d \log M_t}{d \tau_t} &= \underbrace{-\mathcal{E} \left(\frac{\eta}{\eta - \mathcal{E}} \right) \left(\sum_{i=1}^N (w_{i,t}^M - w_{i,t}) \frac{g_i}{\hat{p}_{i,t}} \right)}_{\text{Effect of change in energy composition}} - \underbrace{F \left(\frac{\eta - 1}{\eta} \right) \left(\frac{\eta}{\eta - F} \right) \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}}}_{\text{Effect of change in total energy consumption}} \\
\frac{d \log E_t}{d \tau_t} &= - \underbrace{F \left(\frac{\eta - 1}{\eta} \right) \left(\frac{\eta}{\eta - F} \right) \sum_{i=1}^N w_{i,t} \frac{g_i}{\hat{p}_{i,t}}}_{\text{Effect of change in total energy consumption}}
\end{aligned} \tag{57}$$

The first of which is equation (25) in the main text. Table A.5 computes the elements of this summation for China.

C.2.2 Research subsidy

We now calculate the effect of a subsidy increase for green R&D. It is easier to work with derivatives with respect to $\log(1 - \chi_j)$ in the calculations, so we represent this derivative by d . In the final formulae we use the corresponding units conversion:

$$\frac{d}{d \chi_j} = \frac{-1}{1 - \chi_j} \frac{d}{d \log(1 - \chi_j)}$$

Changes in energy composition. Differentiating equations (35) and (36) and the definition of the post-tax price we get that:

$$\begin{aligned}
d \tilde{e}_i &= -\mathcal{E} d \tilde{p}_i & \forall i \\
d \tilde{p}_i &= d \tilde{p}_i - \frac{g_i \tau}{\hat{p}_i} d \log p_i & \forall i \\
d \tilde{e}_i &= -\eta d \tilde{p}_i - 1 & \forall i \neq j \\
d \log e_j &= 1 - \eta d \log p_j
\end{aligned}$$

Note that unlike in the carbon tax case, the equations differ across i and j . That is because the policy differs explicitly across the different fuels. In the case of a carbon tax, the same expression for the policy governed all the different fuels. In this case that is not true.

We can solve this system of equations to get expression in terms of $\log p_j$. In particular:

$$d \tilde{p}_i = - \left(\frac{1 + J_i \tau d \log p_j}{\eta - \mathcal{E} + J_i \tau} \right)$$

Where $J_i = \mathcal{E} g_i / \hat{p}_i$. For all $i \neq j$, this can be usefully rewritten as

$$\begin{aligned}
d \tilde{p}_i &= C_i(\tau) + D_i(\tau) d \log p_j \\
d \tilde{e}_i &= -\eta C_i(\tau) - \eta D_i(\tau) d \log p_j - 1 \\
d \tilde{p}_i &= \frac{\eta}{\mathcal{E}} C_i(\tau) + \frac{\eta}{\mathcal{E}} D_i(\tau) d \log p_j + \frac{1}{\mathcal{E}}
\end{aligned}$$

Where:

$$C_i(\tau) = \frac{-1}{\eta - \mathcal{E} + J_i\tau} \quad D_i(\tau) = J_i\tau C_i(\tau)$$

Note that:

$$\begin{aligned} C_i(0) &= \frac{-1}{\eta - \mathcal{E}} & D_i(0) &= 0 \\ \lim_{\tau \rightarrow \infty} C_i(\tau) &= \frac{-1}{\eta} & \lim_{\tau \rightarrow \infty} D_i(\tau) &= \frac{-\mathcal{E}}{\eta} \end{aligned}$$

This last line follows from $\lim_{\tau \rightarrow \infty} J_i = \mathcal{E}/\tau$. Note also that the signs here are correct. We are initially working with a proportionate subsidy decrease for good j and so the relative consumption of good i should go up, as it does here. This will get inverted when we do the units conversion at the end.

Changes in energy usage and emissions. Substituting the proportionate changes into equation (55) and solving for $d \log p_j$, we get the general solution for the pre-tax price change:

$$d \log p_j = \frac{B}{A} + \sum_{i \neq j} \frac{M_i}{A}$$

where:

$$\begin{aligned} A &= F - \eta - \eta \sum_{i \neq j} w_i D_i(\tau) + F \sum_{i \neq j} \left(w_P \left(\frac{\eta}{\mathcal{E}} \right) w_i + (1 - w_P) v_i \right) D_i(\tau) + F(1 - w_P) \eta \sum_{i \neq j} (w_i - v_i) D_i(\tau) \\ B &= - \left(w_j + (1 - w_j) F w_P \left(\frac{1}{\mathcal{E}} \right) + F(1 - w_P) (v_j - w_j) \right) \\ M_i &= \left(\left(1 - \frac{F w_P}{\mathcal{E}} - F(1 - w_P) \right) \eta w_i + F(1 - w_P) (\eta - 1) v_i \right) C_i(\tau) \end{aligned}$$

As before, one can now substitute these equations into each other to compute the corresponding general energy quantity changes for general τ . For $\tau = 0$ these again simplify to distinguish the composition and quantity channels, first for the green fuel:

$$\begin{aligned} \frac{d \log p_{j,t}}{d \chi_{j,t}} &= \underbrace{\left(\frac{1 - w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{-1}{\eta - \mathcal{E}} \right)}_{\text{Change in energy mix}} + \underbrace{\left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{-1}{\eta - F} \right)}_{\text{Change in total energy}} \\ \frac{d \log e_{j,t}}{d \chi_{j,t}} &= \underbrace{\left(\frac{1 - w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right)}_{\text{Change in energy mix}} + \underbrace{\left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{F}{\eta - F} \right)}_{\text{Change in total energy}} \end{aligned}$$

Then for the other fuels:

$$\frac{d \log e_{i,t}}{d \chi_{j,t}} = - \left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right) + \left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{F}{\eta - F} \right)$$

And finally, for aggregates:

$$\begin{aligned}\frac{d \log E_t}{d \chi_{j,t}} &= \left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{F}{\eta - F} \right) \\ \frac{d \log M_t}{d \chi_{j,t}} &= \left(\frac{w_{j,t}}{1 - \chi_{j,t}} \right) \left(\frac{F}{\eta - F} - \frac{\mathcal{E}}{\eta - \mathcal{E}} \right)\end{aligned}$$

C.2.3 Production subsidy

The effect of a subsidy for green energy usage is very similar to that for green research. As there, it is easier to work with derivatives with respect to $\log(1 - \delta_j)$ in the calculations, so we represent this derivative by d . In the final formulae we use the corresponding units conversion:

$$\frac{d}{d \delta_j} = \frac{-1}{1 - \delta_j} \frac{d}{d \log(1 - \delta_j)}$$

Changes in energy composition. Differentiating the first-order conditions and the definition of the post-tax price we get that:

$$\begin{aligned}d \tilde{e}_i &= -\mathcal{E} d \tilde{p}_i && \forall i \\ d \tilde{p}_i &= d \tilde{p}_i - 1 - \frac{g_i \tau}{\hat{p}_i} d \log p_i && \forall i \\ d \tilde{e}_i &= -\eta d \tilde{p}_i + 1 && \forall i \neq j \\ d \log e_j &= -\eta d \log p_j - 1\end{aligned}$$

Note that unlike in the carbon tax case, the equations differ across i and j . That is because the policy differs explicitly across the different fuels. In the case of a carbon tax, the same expression for the policy governed all the different fuels. Here, that is not true.

Again, we solve first for $\log p_j$ to get:

$$\begin{aligned}d \tilde{p}_i &= \kappa_i(\tau) + L_i(\tau) d \log p_j \\ d \tilde{e}_i &= 1 - \eta \kappa_i(\tau) - \eta L_i(\tau) d \log p_j \\ d \tilde{p}_i &= -\frac{1}{\mathcal{E}} + \frac{\eta}{\mathcal{E}} \kappa_i(\tau) + \frac{\eta}{\mathcal{E}} L_i(\tau) d \log p_j\end{aligned}$$

Where, as before, $J_i = \mathcal{E} g_i / \hat{p}_i$ and:

$$\kappa_i(\tau) = \frac{1 - \mathcal{E}}{\eta - \mathcal{E} + J_i \tau} \qquad L_i(\tau) = \frac{-J_i \tau}{\eta - \mathcal{E} + J_i \tau}$$

Changes in energy usage and emissions. Substituting the proportionate changes into equation (55) and solving for $d \log p_j$, we get the general solution for the pre-tax price change:

$$d \log p_j = \frac{B'}{A'} + \sum_{i \neq j} \frac{M'_i}{A'}$$

where A', B', M'_i are structurally very similar to the A, B, M_i in the responses to a change in χ_j :

$$A' = F - \eta - \eta \sum_{i \neq j} w_i L_i(\tau) + F \sum_{i \neq j} \left(w_P \left(\frac{\eta}{\mathcal{E}} \right) w_i + (1 - w_P) v_i \right) L_i(\tau) + F(1 - w_P) \eta \sum_{i \neq j} (w_i - v_i) L_i(\tau)$$

$$B' = \left(w_j + (1 - w_j) F w_P \left(\frac{1}{\mathcal{E}} \right) + F(1 - w_P) (v_j - w_j) \right)$$

$$M'_i = \left(\left(1 - \frac{F w_P}{\mathcal{E}} - F(1 - w_P) \right) \eta w_i + F(1 - w_P) (\eta - 1) v_i \right) \kappa_i(\tau)$$

From which changes for general τ can again be easily computed by substituting back into the first order conditions as for the carbon tax. For $\tau = 0$ these once more simplify, first for the green fuel:

$$\frac{d \log p_{j,t}}{d \delta_{j,t}} = - \underbrace{\left(\frac{1 - w_{j,t}}{1 - \delta_{j,t}} \right) \left(\frac{\mathcal{E} - 1}{\eta - \mathcal{E}} \right)}_{\text{Change in relative price}} + \underbrace{\left(\frac{w_{j,t}}{1 - \delta_{j,t}} \right) \left(\frac{1 - F}{\eta - F} \right)}_{\text{Change in total price}}$$

$$\frac{d \log e_{j,t}}{d \delta_{j,t}} = \underbrace{\left(\frac{1 - w_{j,t}}{1 - \delta_{j,t}} \right) (\eta - 1) \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right)}_{\text{Change in energy mix}} + \underbrace{\left(\frac{w_{j,t}}{1 - \delta_{j,t}} \right) (\eta - 1) \left(\frac{F}{\eta - F} \right)}_{\text{Change in total energy}}$$

And for the other fuels:

$$\frac{d \log e_{i,t}}{d \delta_{j,t}} = \left(\frac{-w_{j,t}}{1 - \delta_{j,t}} \right) (\eta - 1) \left(\frac{\mathcal{E}}{\eta - \mathcal{E}} \right) + \left(\frac{w_{j,t}}{1 - \delta_{j,t}} \right) (\eta - 1) \left(\frac{F}{\eta - F} \right)$$

As with the research subsidy, the change in the composition is just the mirror of the response of the green fuel. Total energy use and emissions are thus:

$$\frac{d \log E_t}{d \delta_{j,t}} = w_{j,t} \left(\frac{\eta - 1}{1 - \delta_{j,t}} \right) \left(\frac{F}{\eta - F} \right)$$

$$\frac{d \log M_t}{d \delta_{j,t}} = w_{j,t} \left(\frac{\eta - 1}{1 - \delta_{j,t}} \right) \left(\frac{F}{\eta - F} - \frac{\mathcal{E}}{\eta - \mathcal{E}} \right)$$

The latter of which can be rewritten as equation (26).

From the aggregate responses for the two subsidies it should also be obvious that the contemporaneous responses to the subsidies are in fixed proportion. That is:

$$\frac{d \log E_t}{d \delta_{j,t}} = (\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log E_t}{d \chi_{j,t}}$$

$$\frac{d \log M_t}{d \delta_{j,t}} = (\eta - 1) \left(\frac{1 - \chi_{j,t}}{1 - \delta_{j,t}} \right) \frac{d \log M_t}{d \chi_{j,t}}$$

This is equation (28) in the main text.

C.3 Dynamic derivatives

I now compute the derivatives with respect to policies in t in subsequent periods. These operate via the changes in period t production costs, which affect reference costs in the next period. These in

turn impact production costs in the next period. We can iterate on this to get a full set of marginal impulse responses for all periods after t , via:

$$\frac{dy_{t+s}}{dx_t} = \underbrace{\frac{dy_{t+s}}{d \log \bar{\mathbf{p}}'_{t+s}}}_{N \times 1} \underbrace{\prod_{k=1}^{s-1} \left(\frac{d \log \bar{\mathbf{p}}_{t+k+1}}{d \log \bar{\mathbf{p}}_{t+k}} \right)}_{N \times N} \underbrace{\frac{d \log \bar{\mathbf{p}}_{t+1}}{dx_t}}_{1 \times N} \quad (58)$$

Note that all but the last term are independent of the policy variable x_t . This is important for understanding how policy choices affect outcomes in future. The long-run effects of different policies on *any* outcome variable differ only to the extent that they have different effects on the contemporaneous costs of production. This makes some intuitive sense: in this setting future research builds on past research without distinction between that induced by carbon taxes and that by research subsidies.

Response of reference prices. Differentiating (41) with respect to $\bar{\mathbf{p}}_t$ yields:

$$\frac{d \log \bar{\mathbf{p}}_{t+1}}{d \log \bar{\mathbf{p}}_t} = I + \omega \rho \left(\frac{d \log \mathbf{p}_t}{d \log \bar{\mathbf{p}}_t} - I \right) \quad (59)$$

Derivatives with respect to reference prices. Differentiating equations (47) and (49) w.r.t. \bar{p}_j gives:

$$\begin{aligned} d\tilde{e}_i &= -\mathcal{E}d\tilde{p}_i & \forall i \\ d\tilde{p}_i &= d\tilde{p}_i - \frac{g_i\tau}{\hat{p}_i}d \log p_i & \forall i \\ d\tilde{e}_i &= -\eta d\tilde{p}_i + 1 - \eta & \forall i \neq j \\ d \log e_j &= \eta - 1 - \eta d \log p_j \end{aligned}$$

This is a very similar set of expressions to the R&D subsidy equations, which is somewhat unsurprising as the mechanism at play is similar: the impact of a shift in the research demand functions on equilibrium. An important difference, however, is that the index j now no longer represents a the clean technology, but instead is simply a generic fuel type, whose price has changed in the preceding period.

Solving this in terms of $\log p_j$ gives:

$$d\tilde{p}_i = - \left(\frac{\eta - 1 + J_i\tau d \log p_j}{\eta - \mathcal{E} + J_i\tau} \right)$$

Where $J_i = \mathcal{E}g_i/\hat{p}_i$. For all $i \neq j$, this can be usefully rewritten as

$$\begin{aligned} d\tilde{p}_i &= \hat{C}_i(\tau) + D_i(\tau)d \log p_j \\ d\tilde{e}_i &= -\eta\hat{C}_i(\tau) - \eta D_i(\tau)d \log p_j - \eta + 1 \\ d\tilde{p}_i &= \frac{\eta}{\mathcal{E}}C_i(\tau) + \frac{\eta}{\mathcal{E}}D_i(\tau)d \log p_j + \frac{\eta - 1}{\mathcal{E}} \end{aligned}$$

Where:

$$\hat{C}_i(\tau) = \frac{1 - \eta}{\eta - \mathcal{E} + J_i\tau} = (\eta - 1)C_i(\tau) \qquad D_i(\tau) = \frac{J_i\tau}{\eta - \mathcal{E} + J_i\tau}$$

Substituting these into equation (55) gives the solution for $d \log p_j$:

$$d \log p_j = \frac{\hat{B}}{\hat{A}} + \sum_{i \neq j} \frac{\hat{M}_i}{\hat{A}}$$

where:

$$\begin{aligned} \hat{A} &= F \left(1 - \frac{w_P J_j \tau}{\mathcal{E}} \right) - \eta + \sum_{i \neq j} w_i D_i(\tau) + \\ &\quad F \sum_{i \neq j} \left(w_P \left(\frac{\eta}{\mathcal{E}} \right) w_i + (1 - w_P) v_i \right) D_i(\tau) + F(1 - w_P) \eta \sum_{i \neq j} (w_i - v_i) D_i(\tau) \\ \hat{B} &= -(\eta - 1) \left(w_j + (1 - w_j) F w_P \left(\frac{1}{\mathcal{E}} \right) + F(1 - w_P) (v_j - w_j) \right) \\ \hat{M}_i &= \left(\left(1 - \frac{F w_P}{\mathcal{E}} - F(1 - w_P) \right) \eta w_i + F(1 - w_P) (\eta - 1) v_i \right) \hat{C}_i(\tau) \end{aligned}$$

These are the same formulae as for the R&D subsidy, except for: an extra term pre-multiplying F in \hat{A} , multiplication of B by $\eta - 1$ for \hat{B} , and replacement of C_i by \hat{C}_i for \hat{M}_i . This formula is used in two places. First, to define the derivatives with respect to reference prices, needed to compute $\frac{d \log \mathbf{p}_t}{d \log \bar{\mathbf{p}}_t}$ in equation (59). Second, to give the general formula for the response of relative responses, $d\tilde{p}_i, d\tilde{e}_i, d\tilde{p}_i$, which then give absolute responses for equilibrium variables with respect to the reference price, i.e. the first term in equation (58).

In the special case where $\tau_t = 0$ for all t the dynamic derivatives once more simplify, to:

$$\begin{aligned} \frac{d \log p_{i,t}}{d \log \bar{p}_{i,t}} &= (\eta - 1) \left(\frac{1 - w_{i,t}}{\eta - \mathcal{E}} + \frac{w_{i,t}}{\eta - F} \right) \\ \frac{d \log p_{i,t}}{d \log \bar{p}_{k,t}} &= (\eta - 1) \left(-\frac{w_{k,t}}{\eta - \mathcal{E}} + \frac{w_{k,t}}{\eta - F} \right) \qquad \text{for } i \neq k \end{aligned}$$

Again, the propagation of technological progress breaks down into a composition effect and a quantity effect. When the stock of useful past knowledge about how to produce fuel k increases (that is, $\bar{p}_{k,t}$ declines) there is substitution of research into fuel k and away from fuels $i \neq k$, in proportion to $1/(\eta - \mathcal{E})$. This is the first term in parentheses in the expressions above. In addition, there is a decline in the overall price of energy, increasing usage of all fuels and stimulating energy research in general in proportion to $\eta - F$. This is the second term, common to all fuels.

Computing impulse responses. The preceding is sufficient to compute analytical marginal impulse response with respect to policy x_t , simply substituting in for each of the terms in equation (58), as described above. To calculate the policy decomposition in Figures 7 and A.3, I use a finite difference method, dividing each policy package into a series of small incremental policy changes. I compute the

aggregate response by adding up marginal impulse responses for each small policy change, recomputing weights, prices and the like after each step. An important test of this approach is the difference between the full general equilibrium model solution and the summed partial equilibrium incremental responses. This is very small, validating this approach – see approximation errors in Figures 7 and A.3. Furthermore, given that the marginal partial equilibrium responses are computed independent of the general equilibrium model, this serves as a useful cross-check on the more complicated calculations required to solve the larger model.

C.4 Stability of impulse responses

The impulse responses are an iterative process. This can be seen in equation (58), which shows that the period $t + s$ outcome is the product of three terms: the derivative of reference prices in the initial period; the propagation of reference prices over the following s periods; and then a final term converting the change in reference prices in period $t + s$ to the variable of interest. The key term for stability is thus the middle term, defining the propagation of the reference price. This is not a purely technical point, but essential for how policy affects outcomes. If reference prices in period $t + k + 1$ respond less than one-for-one to reference prices in period $t + k$, then this term decays away and temporary policy changes in period t have only temporary effects. But if this term does not decay to zero then temporary policies can have permanent effects. This is exactly the intuition in Acemoglu et al. (2016), where temporary policies can permanently affect the direction of technical change.

To keep things simple, I focus on the marginal response to the case where $\tau_t = 0$, on the understanding that this creates the largest static and dynamic responses. If stability holds in this case, then it also does in others. I show that the extent of intra-fuel spillovers are essential in determining the stability of the impulse responses, and thus the temporary or permanent nature of responses to policy changes.

The matrix of derivatives of $\log \bar{\mathbf{p}}_{t+1}$ with respect to $\log \bar{\mathbf{p}}_t$ is then:

$$\begin{aligned} \frac{\log \bar{\mathbf{p}}_{t+1}}{\log \bar{\mathbf{p}}_t} &== I + \omega\rho \left(\frac{d \log \mathbf{p}_t}{d \log \bar{\mathbf{p}}_t} - I \right) \\ &= I + \left(\frac{\mathcal{E} - 1}{\eta - \mathcal{E}} \right) \omega\rho + \omega(\eta - 1) \left((\eta - F)^{-1} - (\eta - \mathcal{E})^{-1} \right) \mathcal{W} \\ &\equiv D \end{aligned}$$

Where $\mathcal{W} = \mathbf{1}'\mathbf{w}$ is the matrix of post-tax expenditure weights and D stands for the matrix determining the dynamics of the system. The impulse responses near $\tau_t = 0 \forall t$ are thus:

$$\frac{dy_{t+s}}{dx_t} = \omega \frac{dy_{t+s}}{d \log \bar{\mathbf{p}}_{t+s}} D^{s-1} \rho \frac{d \log \mathbf{p}_{t+1}}{dx_t}$$

The the growth rate of the impulse responses is determined by the eigenvalues of D . In particular, equilibrium dynamics are stable (i.e. the effects a small policy change in period t dissipate as time passes) if $\max \text{eig}(D) < 1$, where $\text{eig}(D)$ is the set of eigenvalues of D . In general, we cannot derive analytical solutions for $\max \text{eig}(D)$. However, we can derive a necessary condition for stability. The

maximum eigenvalue of D is less than unity only if the sum of the eigenvalues is less than the dimension of D , $\sum eig(D) < N$. Because $\sum eig(D) = trace(D)$, and because the trace is additive, we have stability only if:

$$\sum eig(\rho) < \left(\frac{\eta-1}{\mathcal{E}-1}\right) \left(\frac{\mathcal{E}-F}{\eta-F}\right) \in \left[1, \frac{\mathcal{E}(\eta-1)}{\eta(\mathcal{E}-1)}\right] \quad (60)$$

This is where the degree of intra-fuel research spillovers matters. The rows of ρ each sum to unity, the range of the left hand side is $[1, N]$, depending on the degree of linear independence of ρ . If there are no spillovers, then ρ is the identity matrix and so the upper bound N is attained, shrinking the parameter space in which impulse responses are stable. If there are complete spillovers, the rows of ρ are all identical, and the sum of the eigenvalues is 1, all-but guaranteeing stability of the impulse responses.

The intuition for why spillovers matter here is that if there are no spillovers, one fuel can attain an insurmountable technological lead, with a self-reinforcing circle of ever-increasing energy share stimulating ever-higher research driving down costs of production more and more. This is in fact exactly the intuition behind the mechanism in Acemoglu et al. (2016), who show that temporary climate mitigation policies can lead to permanent changes in fuel shares in the long run. With sufficiently large technological spillovers, advances in one fuel drive down costs of production in the others, leading to common movements in production costs and preventing runaway progress in one fuel type.

C.5 Dynamics with symmetric technology diffusion

In the special case where the ρ_{ij} are symmetric, the stability condition simplifies yet further, making the role of intra-fuel spillovers more explicit. That is:

$$\rho_{ij} = \begin{cases} \bar{\rho} + \rho_d & i = j \\ \bar{\rho} & i \neq j \end{cases}$$

Where $\rho_d = 1 - N\bar{\rho}$. In this case, we can decompose ρ into two matrices:

$$\rho = \rho_d I + \bar{\rho} \mathbf{1}_{N \times N}$$

Where I is the identity matrix, and $\mathbf{1}_{N \times N}$ the $N \times N$ matrix of ones. This means that we can decompose D into:

$$D = (1 + a\rho_d)I + (a\bar{\rho}\mathbf{1}_{N \times N} + b\mathcal{W})$$

Adding a multiple of an identity matrix simply increases the eigenvalues by that multiple, so:

$$eig(D) = \{1 + a\rho + \lambda | \lambda \in eig(a\bar{\rho}\mathbf{1}_{N \times N} + b\mathcal{W})\}$$

As $a\bar{\rho}1_{N \times N} + b\mathcal{W}$ is a matrix with identical rows, it only one non-zero eigenvalue, $aN\bar{\rho} + b$. So in this case, the maximum eigenvalue of D is given by:

$$\begin{aligned}
\max \text{eig}(D) &= (1 + a\rho_d) + \max(aN\bar{\rho} + b, 0) \\
&= (1 + a\rho_d) + \max(a(1 - \rho_d) + b, 0) \\
&= \max(1 + a\rho_d, 1 + a + b) \\
&= 1 + \omega \max \left[\left(\frac{\mathcal{E} - 1}{\eta - \mathcal{E}} \right) \rho_d, \frac{F - 1}{\eta - F} \right]
\end{aligned}$$

Note that in any case, the maximum eigenvalue is always greater than one, so there is always growth in the long-run impulse response.

The rate of growth, however, depends on the degree of technological diffusion, parameterized here by ρ_d . With $\mathcal{E} < F$, there is a meaningful threshold for ρ_d . If $\rho_d < (\mathcal{E} - 1)/(F - 1) \times (\eta - \mathcal{E})/(\eta - F)$, then this growth rate is determined by (ω and) the interaction between η and F . If ρ_d is above this threshold, the then the interaction is instead between η and \mathcal{E} . The intuition for this is that if technological innovation diffuses across fuel types (low ρ_d) then the long-run dynamics are determined by an *aggregate* channel – innovation in the energy sector as a whole. But if technological innovation is concentrated (high ρ_d) then long-run dynamics are driven by a *composition* channel – innovation in a specific fuel type.

The difference between the elasticity \mathcal{E} or F and unity defines the strength of the amplifying market size effect. If $F = 1$ then there is no change in energy expenditure, $P_t E_t$, and so no increase in aggregate energy R&D. Likewise, if $\mathcal{E} = 1$ then the *composition* equivalent of this channel fails to operate for similar reasons – fuels expenditure shares are constant. If η gets very large (i.e. returns to R&D decline), then the dynamic effect of R&D is driven to zero.

D Computation

D.1 Basic solution algorithm

Solving for equilibrium within each period involves two iterative processes. The inner process takes a guess of the international oil price, $p_{1,t}$, as given and solves for equilibrium within each region. In each region m , this involves iterating over a guess for energy consumption $E_{m,t}^1$ and checking to see if markets clear. Pseudo-code for solving period t is given below:

Outer Loop

1. Guess the international oil price, $p_{1,t}$. For the first iteration, set $p_{1,t} = p_{1,t-1}$

Inner Loop: in each region m

- i. Guess energy consumption $E_{m,t}$.
 - ii. Invert the CES energy demand curve to find the price $P_{m,t}^d$ at which demand equals $E_{m,t}$
 - iii. Compute energy firms' price $P_{m,t}^s$ which comes from minimizing costs given total energy consumption $E_{m,t}$ and global oil price $p_{1,t}$.
 - iv. Increase (decrease) $E_{m,t}$ if $P_{m,t}^d$ is more (less) than $P_{m,t}^s$
 - v. If $P_{m,t}^d - P_{m,t}^s \neq 0$ return to step ii.
2. Compute global oil demand by summing regional demands.
 3. If global oil demand equals oil supply from region M , stop.
 4. Otherwise, increase (decrease) $p_{1,t}$ if oil demand is more (less) than supply and return to the inner loop.

When both the global oil market and domestic energy markets clear, the period t allocation is determined. The state variables are then updated and the process begins over in period $t + 1$.

D.2 Inverting the model to meet emissions targets

Backing out the policies which will hit a specific emissions target involves using a double bisection algorithm in each period to compute the tax and subsidy levels which deliver the emissions targets and a balanced budget. Given that there are three policies and two targets (emissions and government balance) I solve conditional on the ratio of the research and production subsidies. That is, I start by fixing $\kappa \in [0, 1]$ and in each period solving for the subsidy control $\theta_{m,j,t} \in [0, 1]$, where $\chi_{m,j,t} = \kappa\theta_{m,j,t}$ and $\delta_{m,j,t} = (1 - \kappa)\theta_{m,j,t}$. This allows any ratio of research to production subsidies, with the ratio given by $\kappa/(1 - \kappa)$. However, because government revenue from the carbon tax and expenditure on subsidies depends on equilibrium allocations (higher usage of dirty fuels means more revenue and less expenditure) there is no simple way to compute the budget-balancing subsidy program given the carbon tax. This has to be solved for numerically as well, necessitating the double-bisection approach with the government balance loop nested inside the emissions-target loop. The obvious exception to this is when there is only a carbon tax, in which case the inner loop can be skipped

Outer loop: iterate on the carbon tax to hit the emissions target

1. Guess upper bound and lower on bounds on the carbon tax in each active region, $\tau_{m,t}^{lower}, \tau_{m,t}^{upper}$.
On the the first iteration, set $\tau_{m,t}^{lower} = 0$.
2. Let $\tau_{m,t}^* = \frac{1}{2}(\tau_{m,t}^{lower} + \tau_{m,t}^{upper})$

Inner Loop: iterate on the subsidy control to achieve budget balance

- i. Guess upper bound and lower on bounds on the subsidy control $\theta_{m,t}^{lower}, \theta_{m,t}^{upper}$ in each active region. On the the first iteration, set $\theta_{m,t}^{lower} = 0, \theta_{m,t}^{upper} = 1/\max(\kappa, 1 - \kappa)$.
 - ii. Let $\theta_{m,t}^* = \frac{1}{2}(\theta_{m,t}^{lower} + \theta_{m,t}^{upper})$ for each region.
 - iii. Solve the period t global general equilibrium model with tax and subsidy policies given by $\tau_{m,t}^*, \theta_{m,t}^*$.
 - iv. For each m , update the upper (lower) bound for the subsidy control replacing it with $\theta_{m,t}^*$ if the government balance is in deficit (surplus).
 - v. If $\theta_{m,t}^{lower} - \theta_{m,t}^{upper}$ is small enough stop, else return to step ii.
3. For each m , update the upper (lower) bound for the tax rate control replacing it with $\tau_{m,t}^*$ if region m emissions are below (above) target.
 4. If $\tau_{m,t}^{lower} - \tau_{m,t}^{upper}$ is small enough stop, else return to step 2.

The computationally intensive part of this process is that the double-bisection algorithm already nests the algorithm in Appendix D.1, which is itself a double loop over multiple regions. Even though a little slow, the advantage of using bisection here is twofold: it is robust – reliably converging in a wide range of use cases – and it is guaranteed to find a solution of given precision within a fixed time. Both tax and subsidy policies solve to acceptable tolerances within a dozen iterations, putting a hard cap of 144 solutions in each period.⁴¹ With almost forty periods to solve for this remains a lengthy process, but because constant savings rates mean we can solve period-by-period, it is a manageable one. As a result, for two regions setting emissions targets this process converges overnight on a modern laptop. If savings rates were endogenous, though, this would need to be iterated on dozens of times; computing solutions with multiple parameter values, as I do in the main text, would take months. Moreover, because the processes here are inherently sequential – each bisection relies on the previous iteration, and each period on solving the one before it – there is very little scope to speed the process through parallelization.

⁴¹Bisection shrinks the error by half with each step, implying a 12-step error of $0.5^{12} \simeq 0.025\%$ of the initial interval. Because subsidies are constrained to the unit interval, with means an error on the subsidy rate of no more than 0.025 percentage points. If the initial upper bound on taxes is \$1000, then this means an error on the tax rate of around 25 cents.

D.3 Solving the model with non-constant savings rate

With non-constant savings rate, the Euler equation is replaced with equation (31), reproduced here:

$$\left(\frac{C_{m,t+1}}{C_{m,t}}\right)^\gamma = \beta (i_{m,t+1} + 1 - \delta_K) \quad (61)$$

Suppressing dependence on m and substituting in for the savings rate, rental rate for capital, and households' period budget constraint, we get that:

$$\begin{aligned} s_t &= \frac{1}{1 + \Gamma_t} \left(\frac{\alpha(1 + g_{t+1})}{\beta^{-1}((1 + g_{t+1})(1 - s_{t+1})/(1 - s_t))^\gamma - (1 - \delta_K)} - (1 - \delta)k_t \right) \\ &\equiv f(s_t; s_{t+1}, k_t, g_{t+1}, \Gamma_t) \end{aligned}$$

Where $(1 + g_{t+1}) = \hat{Y}_{t+1}/\hat{Y}_t$ is the growth in earned income and $k_t = K_t/\hat{Y}_t$ is the capital to (earned) income ratio. Note that given current capital ratio k , income growth g' , unearned income ratio Γ , and next period savings s' , the savings rate s which solves the Euler equation is the solution to a fixed point problem:

$$s = f(s; s', k, g', \Gamma)$$

Note also that on a long-run balanced growth path where income growth, the capital ratio, and lump sum transfers are constant at $\bar{g}, \bar{k}, \bar{\Gamma}$ respectively, the long-run savings rate \bar{s} is given by:

$$\begin{aligned} \bar{s} &= \frac{\alpha\beta(1 + \bar{g})(\bar{g} + \delta_K)}{(1 + \bar{g} + \bar{\Gamma})(\bar{g} + \delta_K)((1 + \bar{g})^\gamma - \beta(1 - \delta_K))} \\ &\equiv h(\bar{k}, \bar{g}, \bar{\Gamma}) \end{aligned}$$

Note that when $\gamma = 1, \delta_K = 1, \Gamma = 0$, this immediately simplifies to $\bar{s} = \alpha\beta$, the savings rate in the full-depreciation version of the model.

I use these to compute the general equilibrium solution with a the savings rate which satisfies the Euler equation. The basic idea is to assume that by the end of the simulation that the model has reached a balanced growth path in each region. Given that the simulation period is nearly two centuries this seem reasonable. I then simulate the model for a given path of savings rates, and then use the terminal period to define the long-run values $\bar{g}, \bar{k}, \bar{\Gamma}$ in each region. From this, I compute a terminal period savings rate s_T using the long-run formula $h(\cdot)$, which I use to initiate a backward recursion, using the values from the simulation for k_t, g_{t+1}, Γ_t in $f(\cdot)$ to compute the sequence of savings rates consistent with the Euler equation back to period 1. If this is sufficiently close to the assumed sequence of savings rates using in the simulation then we are done. If not, I update the savings rate and try again until convergence. Psuedo-code is below.

Part 1: Check the Euler equation errors on the current guess

1. Guess the sequence of savings rates for all regions, $s_{m,t}$.
2. Solve the general equilibrium model fixing the savings rates at these guesses. This is a minor

modification to the algorithm in Appendix D.1, changing the rule for updating the states from one period to the next by replacing the savings rate $\alpha\beta$ with the appropriate entry from this sequence instead.

3. Use the model values for capital, income growth, and unearned income share to compute the Euler equation errors in every region and period, using $error_{m,t} = s_{m,t} - f(s_{m,t}, s_{m,t+1}, k_{m,t}, g_{m,t+1}, \Gamma_{m,t})$
4. If the Euler equation errors are small enough, stop. Otherwise, continue.

Part 2: Update the savings rate guess

5. Compute a new guess for terminal savings in each region using $s_{m,T}^* = h(k_{m,T}, g_{m,T}, \Gamma_{m,T})$.
6. Iterate backwards to update the savings rate to $s_{m,t}^*$ using the Euler equation and simulated values for capital, income growth, and unearned income share, solving in each period and region the fixed point problem

$$s_{m,t}^* = f(s_{m,t}^*, s_{m,t+1}^*, k_{m,t}, g_{m,t+1}, \Gamma_{m,t})$$

NB: This fixed point problem is solved via an iterative loop of its own, but is relatively straightforward to implement, so not discussed further here.

7. Return to step 1 using the updated guess for the savings rate, replacing $s_{m,t}$ with $s_{m,t}^*$.