To everything there is a season: Carbon pricing, research subsidies, and the transition to fossil-free energy☆

Rob Hart[a]

[a]*Department of Economics, SLU, Box 7013, 750 07 Uppsala, Sweden.*

Abstract

We develop a climate policy model with directed technological change (DTC) in the energy sector. The model delivers both analytical and numerical results which give a clear understanding of the respective roles of research subsidies and emissions pricing. By contrast to existing models with DTC, ours is close in structure to recent integrated assessment models, leading to dramatically different results. Although clean-research subsidies are substantial initially, they subsequently decline whereas emissions taxes increase without bound. Furthermore, emissions taxes are far more important than research subsidies: in our baseline parameterization, a regulator unable to tax can only achieve 36 percent of potential benefits, whereas a regulator unable to subsidize can achieve 91 percent of potential benefits.

JEL: O11, O33, Q54, Q55.

Keywords:
Directed technological change, Knowledge spillovers, Energy, Climate change.

☆Thanks to the Swedish Energy Agency (project 37758-1) for financial support. Thanks also to the following for valuable input: Stefan Ambec, Anton Bondarev, Carolyn Fischer, Peter Kjær Kruse-Andersen, Francesco Ricci, Sjak Smulders, Andreas Stephan, several anonymous referees, and seminar participants at EAERE Gothenburg, the LSE conference on the economics of innovation, diffusion, growth and the environment, NTNU (Trondheim), RATIO, the Swedish national conference in economics in Gothenburg, SURED 2016, and Toulouse School of Economics.

Email address: `rob.hart@slu.se` (Rob Hart)

1. Introduction

This paper adds to the climate-policy literature regarding the relationship between carbon pricing and subsidies to clean-energy research. Until 2012 the consensus in this literature was that emissions pricing is crucial while the role of research subsidies is limited; for instance, in an influential paper Fischer and Newell (2008) perform a ranking of instruments and put emissions pricing top and R&D subsidies sixth and last. However, Acemoglu et al. (2012) develop a model with directed technological change (henceforth *DTC*) and conclude that (p.159) 'optimal regulation involves small carbon taxes because research subsidies are able to redirect innovation to clean technologies before there is extensive environmental damage', and Acemoglu et al. (2016) build a related model in which they use microdata for the modelling of competition in production and innovation, and argue (p.55–56) that 'Research subsidies are powerful in redirecting technological change, and given this, it is not worth distorting the initial production too much by introducing very high carbon taxes . . . the social planner relies heavily on research subsidies.' Finally, Greaker et al. (2018) back up the conclusions of Acemoglu et al. (2012) using an extended model. In this paper we develop a model with energy-sector DTC with a much more realistic structure than that used in existing literature, and use this model to deliver both analytical and numerical results which give a better understanding of the respective roles of emissions pricing and research subsidies, and restore the primacy of emissions pricing in the ranking of instruments.

In more detail, the contributions of this paper are as follows. The first contribution is to build a model which includes both DTC and an overall structure similar to IAMs such as Golosov et al. (2014). We thus have a structure in which inputs into the final good sector are labour and energy services such as electricity and motive power, energy services are produced using either fossil fuels or renewables, and emissions flow in proportion to the burning of fossil fuels. This is important because the structure determines the available abatement and regulatory options. In our structure there are three mechanisms for reducing carbon emissions: substitute labour for energy services in the final-good sector; raise fossil-augmenting knowledge in the production of energy services; and substitute from fossil inputs to clean inputs in the production of energy services. By contrast, in Acemoglu et al. (2012) and Greaker et al. (2018) emissions are in proportion to the production of dirty energy intermediates rather than the burning of fossil fuels, and the only way to reduce emissions is to switch from dirty to clean intermediates (it is not possible to economize on fossil fuels in any conventional sense).[1]

[1]A consequence of our structure is that we must account for physical limits on the productivity of primary-energy inputs (such as coal or wind) in delivering energy services (such as electricity or motive power); a quantity of fossil fuel containing 1 MW of chemical energy can under no circumstances deliver more than 1 MW of electricity to the final-good sector. Such limits are rarely acknowledged in the literature on technological change, growth, and energy policy — an exception being Lemoine (2015) — but they have important

The second contribution is to deliver, in a first-best regulated economy, analytical results that clarify the roles of emissions prices and research subsidies. We highlight two of these analytical results here: firstly, given a damage function following Golosov et al. (2014) and Acemoglu et al. (2016) the Pigovian tax in first best grows without bound (because marginal damages track potential GDP) hence a transition from fossil to clean energy is inevitable in first best as long as the unit cost of clean energy is bounded above;[2] and secondly, subsidies to clean-energy research should be higher than subsidies to fossil-energy research during a transition from a state in which fossil energy dominates to a state in which clean energy dominates, the reason being that during such a transition each clean-energy researcher has more future researchers standing on their shoulders than each fossil researcher. Hart (2008) notes (Proposition 4) that a high rate of investment in a given sector is linked to a high rate of spillovers and hence more severe undersupply of research in that sector, but Heggedal (2015) is the first to thoroughly investigate and highlight the need for higher subsidies in growing sectors; see also Greaker et al. (2018).

The third contribution is to develop a model of DTC with two novel elements which make crucial differences to the properties of the model. (1) Total research inputs into the energy sector are endogenous, so the allocation problem in the model is not simply where to put researchers of whom there is a fixed number, the problem is how many researchers to put into clean and fossil research respectively. The result is that the regulator typically subsidizes both clean and fossil research.[3] (2) We allow for knowledge spillovers between three knowledge types, labour-augmenting, fossil-augmenting, and clean-energy-augmenting knowledge. Clean–clean and fossil–fossil spillovers imply path dependence, but intersectoral spillovers (ruled out by construction in Acemoglu et al. (2012) and Greaker et al. (2018)) weaken path-dependence because they imply that the productivity of research in a laggard sector is boosted by knowledge in the other sectors.[4]

consequences, implying that the price of energy services should approach a lower limit over time rather than approaching zero over time.

[2]The results of Acemoglu et al. (2016) are mystifying in the light of this theoretical result. In Acemoglu et al. marginal damages are independent of the atmospheric concentration, so the social cost of carbon should be proportional to Y/δ, where Y is current gross production and δ is a discount rate equal to the interest rate plus the decay rate of carbon minus the growth rate. As long as Y keeps growing and the discount rate doesn't grow, the emissions tax should keep growing. But in Acemoglu et al. it follows a hump-shaped path in all simulations.

[3]In Acemoglu et al. (2012), Acemoglu et al. (2016), and Greaker et al. (2018) the question facing the regulator is essentially 'How high do clean-research subsidies need to be in order to ensure that all of the energy-researchers in the economy switch from fossil-augmenting research to clean-augmenting research?'

[4]In Acemoglu et al. (2012) there are no spillovers between knowledge types, hence the degree of path dependence and the strength of lock-in effects are exaggerated, since catching up is more difficult than in reality. This is also standard in many other models of energy and environment, including Smulders and de Nooij (2003), Gerlagh (2008), Fischer and Newell (2008), and Hassler et al. (2012). In Acemoglu et al. (2016) there are spillovers from fossil to clean knowledge in each of a continuum of sectors, modelled as a small probability that there will be a quantum leap in clean productivity up to the level of fossil productivity. Outside the climate/energy literature, Acemoglu (2002) develops and analyses a model of DTC with spillovers between two sectors, and we build on this model to a degree; see also Nordhaus (1973) and Hart (2013), who discuss the shortcomings of the model without spillovers, and what is needed to replace it.

The fourth contribution is to take 'second-best' seriously. That is—following the terminology of Lipsey and Lancaster (1956)—we add constraints which prevent the achievement of the Paretian optimum, and investigate the effect on the optimal choices of all the instruments. If we were to adopt a standard framework in which a Paretian (first-best) allocation could be achieved through regulatory instruments, research subsidies would be very high to both clean and fossil-energy research, because all forms of research are severely undersupplied in such a framework (as shown by Jones and Williams, 1998). To avoid this uninteresting result and ensure a more relevant model we introduce deadweight losses caused by research subsidization and allow for the 'stepping on toes' effect described by Jones and Williams (1998), and find the unique parameterization of these effects which accounts for the observed aggregate level of subsidization to all forms of research in the US. We also introduce a deadweight loss from emissions taxation which is in proportion to the costs imposed by the tax, which is consistent with the observation that levels of emissions taxation currently observed in most jurisdictions are below the optimal level suggested by researchers such as Nordhaus (2008) and Golosov et al. (2014).

The novel features described above allow us to calibrate the model such that the starting point is broadly consistent with historical data, in the sense that if we started the model in 1904 instead of 2004 then the model economy would develop in such a way that demand for fossil and clean energy would be broadly consistent with observations, while the state of the model in 2004 would be broadly consistent with our chosen starting point.[5] We take this parameterization and perform two policy experiments, and sensitivity analysis.

In the first policy experiment, we compare three regulatory options: (i) laissez-faire; (ii) optimal regulation (in second-best); and (iii) first-best regulation in the hypothetical absence of deadweight losses. We find that in the regulated economy (in second-best) both emissions taxes and subsidies to clean and fossil-energy research are used, and there is a dramatic reduction in carbon emissions compared to laissez-faire. Subsidies to clean-energy research are initially large, and gradually decline as the transition to clean energy progresses. Subsidies to fossil-energy research are modest, which is partly because when fossil energy is cheaper this makes the transition to clean energy harder to achieve. The emissions tax starts low and increases monotonically, in line with what we expect from the analytical results. However, it is initially well below the marginal external cost of emissions (MEC), because initially—when the tax only makes up a small part of the cost of fossil inputs—the elasticity of emissions to the tax is low, while the tax causes significant deadweight

[5] This may be clarified through a counterexample. Acemoglu et al. (2012) set up a model with very strong path dependence in clean and dirty research, and choose (in their baseline case) a starting point in which the productivities of clean and fossil inputs are almost equal. This is not consistent with historical evidence since given strong path dependence and a history (before the first period of the model) of fossil dominance, fossil knowledge should be far ahead of clean in the first period.

losses. These results contrast starkly with those of Acemoglu et al. (2012, 2016), where a very high initial subsidy to clean-energy research rapidly declines to zero, subsidies to fossil-augmenting research are zero, and emissions taxes are low and declining (Acemoglu et al., 2012) or hump-shaped (Acemoglu et al., 2016).

In our second policy experiment we investigate the relative importance of the emissions tax and the research subsidies, and the losses associated with climate-change denial by the regulator. We find that climate policy is very important, and that the emissions tax is far more important than subsidies. When a regulator can only use the tax (and cannot subsidize research), 91 percent of the utility gain from second-best optimal policy can still be achieved; when the regulator can only use research subsidies (and cannot tax emissions), only 36 percent of utility gains can be achieved; and optimal policy given climate-change denial yields just 15 percent of the benefits of truly optimal policy. This is radically different to Greaker et al. (2018), who find that the tax is of marginal importance whereas the research subsidy is indispensable. Greaker et al.'s result is a straightforward consequence of the construction of their model, where the regulatory problem (as in the models of Acemoglu et al., 2012, 2016) is effectively how to ensure that all of the fixed number of researchers switch to clean research; given such a switch, the clean intermediates will soon be cheaper than fossil-based intermediates, and the climate problem is solved.

Sensitivity analysis shows that the conclusions are robust in the sense that changes in parameters lead to predictable and limited changes in the results: the key to the results is the construction of the model rather than the exact parameterization.

Finally note two relevant papers on related topics. First, Peretto (2008) develops a model of DTC with a structure related to this one, with firms producing differentiated goods and performing in-house research. Second, the climate policy model developed by Fried (2018) has many things in common with this one, including a nested structure with an energy intermediate produced using either fossil or clean technology, and spillovers between knowledge types which limit the strength of path dependence. Fried's aim is to evaluate the quantitative impact of a carbon tax on green technologies, which she finds to be large, a result which is consistent with our result that a climate tax alone can achieve over 90 percent of the benefits of an optimal combination of a tax and research subsidies.

The paper consists of three main sections. In Section 2 we develop the model, paying particular attention to DTC but otherwise specifying only as much as we need to derive our analytical results. We set out the options open to the regulator and explain the presence of deadweight losses. In Section 3 we derive first-order conditions necessary for a second-best optimal solution in the market economy. We then perform a similar exercise for a benevolent social planner (who does not need to use economic instruments and therefore does not have to worry about deadweight losses), and compare the two to derive analytical results regarding optimal regulation in first-best. In Section 4 we turn to optimal regulation in second best, i.e. when there are deadweight losses associated with the use of economic instruments. To derive these results we must specify the model fully, parameterize, and simulate numerically. Section 5 concludes.

2. The model

2.1. The environment

In the model economy there is a fixed number of households that supply labour L inelastically. Their utility U is a discounted function of net consumption X_t^* over infinite periods, with constant marginal utility of consumption, and discount factor per period β. Net consumption X_t^* is gross consumption X_t modified by a pollution damage factor $\exp(-\gamma_1 S_t^{\gamma_2})$, where $\gamma_1 > 0$, $\gamma_2 \geq 1$, and S_t is the stock of pollution. Pollution can thus be interpreted either as affecting utility directly, or as affecting all production sectors equally and therefore leading to a reduction in net consumption without affecting relative prices or quantities of traded goods:

$$U = \sum_{t=0}^{\infty} \beta^t X_t \exp(-\gamma_1 S_t^{\gamma_2}). \tag{1}$$

Final-good production is competitive, and there is a representative final-good production firm with the production function

$$Y_t = (A_{Lt} L)^{1-\alpha} R_t^{\alpha}, \tag{2}$$

where A_{Lt} is labour productivity, R_t is energy services, α is a parameter equal to the factor share of energy in total production Y. The price of the final good is normalized to 1, and labour productivity grows exogenously by a factor $1+\theta$ each period. Energy services R should be thought of as (for instance) electricity and motive power, i.e. things that are used directly in the production of final goods such as transport and communication.[6]

There is a unit continuum of firms producing energy services. The firms—indexed by i—buy their inputs on competitive markets, and compete monopolistically in selling their products, quantities R_i, prices p_{ri}, and in symmetric equilibrium the elasticity of demand for

[6] Note that we abstract from capital and set the interest rate exogenously. Since we assume perfect information and a constant growth rate of labour productivity A_L, if we endogenized the interest rate based on the Euler equation and CRRA utility, and included neoclassical capital accumulation, we would obtain an almost constant interest rate, and the quantity of capital would track growth in A_L; if we write the resultant production function as $(A_L L)^{1-\alpha-\beta} K^{\beta} R^{\alpha} = (A_L L)^{1-\alpha} [K/(A_L L)]^{\beta} R^{\alpha}$ it should be clear that this implies that the addition of capital would make no significant difference to the behaviour of the model economy.

each firm's product is η:

$$R = \left[\int_0^1 R_i^\eta \mathrm{d}i\right]^{1/\eta}.$$

In practice, market power may arise in the energy sector due to barriers to entry, fixed costs and scale economies, etc. In the model, some form of market power is necessary in order to provide an incentive for research, as is standard in the endogenous growth literature. We set the mass of firms and the elasticity of substitution between inputs exogenously.[7]

Production of R requires primary energy inputs C (clean) and D (dirty, i.e. fossil). When purely clean inputs are used, the production function of firm i is $R_{it} = A_{cit}X_{cit}$, where X_{ci} is the quantity of final goods devoted to the production of clean energy services, and A_{ci} is an index of the productivity of these goods in both capturing clean energy (for instance from the wind or sunlight) and converting it into energy services. On the other hand, when purely fossil inputs are used by firm i, we have $R_{it} = A_{dit}D_{it}$, where D_{it} is the flow of fossil inputs used at t, which is equal to the flow of extraction, $D_{it} = A_{dx}X_{dit}$, where A_{dx} is the productivity of extraction inputs and X_{di} is their quantity. We assume that the extraction input is the final good, hence the price is constant (normalized to 1) and productivity is also constant since the final good is unchanging. This implies that unit extraction costs p_d are constant and equal to $1/A_{dx}$. Finally, energy services from fossil and clean sources are imperfect substitutes for one another, and when both sources are used we have

$$R_{it} = [(A_{cit}X_{cit})^\varepsilon + (A_{dit}D_{it})^\varepsilon]^{1/\varepsilon}, \tag{3}$$

so the elasticity of substitution between clean and fossil energy is $1/(1-\varepsilon)$; $\varepsilon \in (0,1)$.[8]

Burning fossil fuels D leads to emissions P where $P = D$, and emissions add to a stock of pollution S which evolves as follows:

$$S_t = G(H_t) + P_t, \tag{4}$$

where H_t is the set of all emissions levels P for periods prior to t—so $H_t = (P_{t-1}, P_{t-2}, \ldots)$

[7] For the parameterization of η see page 20, and for further discussion on the market structure see Appendix A.2.

[8] Here we briefly compare and contrast our energy-sector model with that of Golosov et al. (2014). Firstly, we simplify by considering just one fossil fuel with constant extraction costs; Golosov et al. have two fuels, oil which is free to extract but scarce, and coal which has constant extraction costs but not scarce. We simplify in order to focus on the key issue of this paper, i.e. the respective roles of research subsidies and emissions pricing in climate policy; it would be interesting to test an extension to two fossil inputs as in Golosov et al. Secondly, our extraction input is the final good (which has constant productivity and constant price) whereas theirs is labour (which has increasing productivity and increasing price). This difference is of 'academic' interest, since the result is the same in each case: constant unit extraction costs. For more on why the extraction costs and prices of non-renewable natural resources tend to be approximately constant see Hart (2016).

—and G is a strictly increasing function of each of its arguments. There is a limited initial stock of the fossil resource, Q_0, which is homogeneous. Thus we have

$$\sum_{t=0}^{\infty} D_t \leq Q_0. \tag{5}$$

Finally, the factor productivities of each firm in the energy-services sector are functions of in-house research by that firm, building on general knowledge both within the energy sector and in the economy as a whole. For details see below.

2.2. Knowledge growth and productivity

We model productivity as a function of knowledge, but distinct from it. It is common in the growth literature to assume that knowledge and productivity are synonymous, and that since there is no obvious limit on knowledge there is no limit on productivity. However, this approach is not reasonable when applied to fossil fuels, where the key step is the conversion of chemical energy in fossil fuels to electrical or kinetic energy to power the economy: there is a limit on the energy we can get out of a given quantity of fossil fuel, and when this limit is approached we say that the efficiency of the process approaches 100 percent while the productivity approaches the limit. We assume that this limit is approached as knowledge approaches infinity. The distinction between knowledge and productivity is not just a theoretical curiosity, since opportunities for increasing the efficiency with which we can obtain either motive power or electricity from fossil fuels are almost exhausted; for instance, efficiency of the best modern coal-fired power stations is at approximately 75 percent of the thermodynamic limit.[9] Note that there is no limit on the efficiency with which we can use energy services to generate utility in the final-good production function.

Consider the generation of motive power from fossil fuels. Define the productivity of the process in firm i as A_{di} (as in equation 3) and define the theoretical maximum productivity as $\bar{A}_d$, which has the same units as A_{di}. Assuming a similar situation in the clean sector we define

$$A_{cit} = \bar{A}_c \frac{K_{cit}}{K^* + K_{cit}} \quad \text{and} \quad A_{dit} = \bar{A}_d \frac{K_{dit}}{K^* + K_{dit}}, \tag{6}$$

so when knowledge is zero productivity is zero, and when knowledge approaches infinity productivity approaches the limit, and the positive parameter K^* is used to calibrate K in

[9]The best modern coal-fired power stations reach around 45 percent thermal efficiency (45 percent of the energy released on burning is converted to electricity), whereas the theoretical limit based on the Carnot cycle is around 60 percent. See Holman (1980) for theoretical background, and IEA (2012) for a recent discussion.

relation to A. Define K_{ct} and K_{dt} as the aggregate knowledge stocks, and assume that

$$K_{ct} = \int_0^1 K_{cit} \mathrm{d}i \quad \text{and} \quad K_{dt} = \int_0^1 K_{dit} \mathrm{d}i. \tag{7}$$

So in symmetric equilibrium the levels of aggregate knowledge are simply equal to the knowledge levels of the representative energy firm in the same period. And in symmetric equilibrium we have

$$A_{ct} = \bar{A}_c \frac{K_{ct}}{K^* + K_{ct}} \quad \text{and} \quad A_{dt} = \bar{A}_d \frac{K_{dt}}{K^* + K_{dt}}, \tag{8}$$

where A_{ct} and A_{dt} are productivities of the representative firm.

We now turn to the knowledge production functions, that is the functions linking a firm's in-house research effort to the subsequent knowledge within the firm, which in turn determines the productivity of the firm's inputs through equation 6. The knowledge production functions are as follows:

$$K_{cit} = (1-\delta_K)K_{ct-1} + A_{Lt-1}^{\sigma_1} \left(K_{ct-1} + \sigma_2 K_{dt-1}\right)^{1-\sigma_1} \zeta Z_{cit}^{\phi}, \tag{9}$$

$$\text{and} \qquad K_{dit} = (1-\delta_K)K_{dt-1} + A_{Lt-1}^{\sigma_1} \left(K_{dt-1} + \sigma_2 K_{ct-1}\right)^{1-\sigma_1} \zeta Z_{dit}^{\phi}, \tag{10}$$

where $\phi \in (0,1]$, $\delta_K \in (0,1]$, $\sigma_1 \in (0,1)$, $\sigma_2 \in (0,1)$, and Z_{cit} and Z_{dit} are the quantities of the research input invested by firm i in the respective sectors. The research input Z is supplied at price w_z. We do not specify at this stage what the input is, since it is not necessary to do so in order to derive the analytical results of Section 3, but note that in the fully specified model of Section 4 the input is labour.

The functions imply that firm i's knowledge in period $t+1$ builds purely on general knowledge in the economy, and not at all on that firm's knowledge in period t. Thus we simplify the process by which the benefits of private investment in knowledge filter out into the public domain, treating this gradual process as a discrete step; in the parameterized model we set the period length to 10 years.[10] In the absence of investment, knowledge stocks decay, but investment adds to the knowledge stock, building on existing knowledge. Furthermore, since σ_1 and σ_2 are both strictly positive this implies that research productivity in a given sector (for instance clean energy) is an increasing function not only of existing

[10]The interpretation is as follows. Normally, in continuous time, we expect the effect of a past investment decision on a firm's current knowledge stock to attenuate gradually over time, as the old knowledge becomes outdated and irrelevant, and new knowledge builds on general knowledge in the overall economy. We make this process discrete; a period in the model is thus equal to the lifetime of an investment, and all firms invest simultaneously. This feature simplifies the dynamics of the model since investing firms—each period—face an essentially static problem. Note that this way of modelling is not uncommon in the literature: Aghion and Howitt (1992) use essentially the same assumption, although there the length of periods is uncertain and there is only one leading firm, while Acemoglu et al. (2012), in their model of DTC and fossil-fuel demand, assume that patents last exactly one period in a discrete-time model, giving researchers a static problem.

clean knowledge, but also of knowledge in the other sectors (in this case, fossil energy and labour). The following thought experiment shows why this should be the case. Consider an imaginary global economy run purely on wind power up to the year 1900, when windmills made of wood with cloth sails are used to generate motive power. At this point a workable coal-fired steam engine is developed, fossil power takes over, and the windmills are abandoned. Over the next 100 years a myriad of technologies (including computers and electric turbines) are developed, and fossil fuels power the economy. In the year 2000 negative effects from fossil fuels are discovered, and wind power is revived; clearly the knowledge developed in the intervening period—both in the overall economy and in the fossil-energy sector—will be a major help to wind-power researchers. Finally, the parameter ϕ—the elasticity of new knowledge production to research inputs—allows for the 'stepping on toes' effect mentioned in the introduction, the basic idea of which is that extra researchers duplicate the discoveries of existing research.

2.3. Regulation and deadweight losses

In the market economy there is a benevolent regulator who uses economic instruments to affect the allocation chosen by agents in order to maximize U. The regulator has four instruments available to her: a subsidy to the energy aggregate R, such that the price paid for this good is $p_r(1-s_r)$; an emissions tax τ_d on each unit of the fossil input D, such that the price paid for fossil fuels is $p_d+\tau_d$; and subsidies to clean and fossil research such that their prices are $w_z(1-s_{cz})$ and $w_z(1-s_{dz})$ respectively. Net payments are returned lump-sum, and we assume that there is no cost of public funds. All of the instruments are constrained to be non-negative. Thus the regulator can affect the price of energy and hence the quantity demanded by the final-good producers, the price of fossil inputs used by the representative energy producer, and investment in C-augmenting and D-augmenting knowledge by the representative energy producer.

Recall from the introduction (page 3) that given endogenous total research effort in the energy sector we would expect very large subsidies to both kinds of research, since research is severely undersupplied in the market compared to first best. Since such high subsidies are not typically observed in any sector, there must be something else going on not captured by standard models. And we need to introduce 'something else' into our model to avoid this uninteresting (and unrealistic) result. Furthermore, since we adapted the model to account for low observed research subsidies, we also adapt the model to account for the low observed carbon price even in 'active' coalitions such as the EU. In both cases we introduce deadweight losses as a result of the instruments.

In the research sector we assume that subsidy payments cause firms to be less efficient in their use of the research input than they would be if they faced optimal prices (i.e. the full price of the input and a shadow price of knowledge equal to the regulator's shadow price).[11] We thus distinguish between effective research inputs, Z_c, and total research inputs which we define as $Z_c(1+\omega_c)$ (with equivalent definitions in the fossil sector). These extra (unproductive) research inputs are in proportion to the square of the subsidy:

$$\text{Total research inputs (clean)} = Z_c(1+\omega_c) = Z_c(1+\delta_z s_{zc}^2),$$

where δ_z is a positive parameter. In the parameterized model we find the unique value of δ_z which (together with ϕ, the elasticity of new knowledge to research) accounts for the observed aggregate level of subsidization to all forms of research in the US (see Section 4.2). We find that $\delta_z = 1.14$, hence given a 30 percent subsidy, 9 percent of research inputs into the sector are wasted. And if the regulator pays (for instance) 60 percent of research costs then 29 percent of research inputs are wasted.[12]

The Pigovian tax on carbon emissions, τ_d, is paid by energy producers, and raises the unit cost of the energy intermediate. In an economy with international trade and climate negotiations, taxation in one country or coalition of countries may cause 'carbon leakage' since it raises the price of energy intermediates within the coalition and may therefore cause energy-intensive industries to move abroad, reducing the benefit of the tax while increasing costs (see for instance Di Maria and van der Werf, 2008). This may explain the reluctance of coalitions such as the EU to impose significant general carbon prices (for instance by ensuring a higher permit price in the EU ETS) despite taking many other costly measures to reduce their carbon emissions.

It is beyond the scope of this paper to include international trade and climate negotiations. However, we wish to test the idea that research subsidies may be an effective substitute for carbon pricing when a Pigovian carbon price causes other market distortions. To do so we introduce such distortions in the theoretical model in a rudimentary way, and compare results when the distortion is present with results when it is absent; the latter scenario can be interpreted as yielding optimal climate policy in a global economy in which there is full coordination.

We assume that the deadweight costs of the tax are simply an enforcement cost—due to the cost to the regulator of preventing tax avoidance—and that these costs are in proportion

[11] Raising the shadow price of knowledge would require the regulator to correctly judge the value of each discovery made by each firm, an impossible task. So what a regulator typically does is to reduce to cost of research inputs instead. The idea behind the model is that the bigger these subsidies are, the more likely it is that unproductive researchers will apply for them, and the higher the likelihood that unproductive research will be funded. As pointed out by a referee, a full model of this effect might follow Jaimovich and Rebelo (2017) in assuming heterogeneous research labour inputs. See Arrow (1962) for a classic discussion of the case for research subsidies, and Klette et al. (2000) for an empirical analysis.

[12] Total research costs for firm i in the clean sector are $w_{zt}Z_{cit}(1+\delta_z s_{zct}^2)(1-s_{zct})$, as can be seen in equation 12 below. And the expression for the percentage of inputs wasted is $\delta_z s_{zc}^2/(1+\delta_z s_{zc}^2)$. Again, equivalent expressions apply in the fossil sector. Note that firms still benefit from the subsidy when $\delta_z > 0$, but not as much as they would do if they could perfectly control their costs.

to the potential savings that firms could make by avoiding the tax. We define this potential saving as the difference between the unit cost of the energy intermediate when paying the tax and when not paying the tax, multiplied by total production. Defining these costs as c_{rt} (paying the tax) and c^*_{rt} (avoiding the tax) we have

$$\mathrm{DWL}_D = (c_{rt} - c^*_{rt})\delta_\tau R_t,$$

where δ_τ is a positive parameter. The equations for c_{rt} and c^*_{rt} (equation C.5) are derived and presented in Appendix C.1.[13]

3. Market equilibrium, and optimal regulation in first-best

3.1. The market solution

We now define a market equilibrium of the above economy, given the regulator's choice of $[s_{rt}, \tau_{dt}, s_{zct}, s_{zdt}]_{t=0}^{\infty}$, and assuming that cumulative extraction of fossil fuels approaches a limit below Q_0, so the scarcity rent is zero. To do so we set up the profit-maximization problem of the representative final-good producer in period t, normalizing the price of the final good to 1,

$$\max_{L_t, R_t} \pi_t = (A_{Lt}L_t)^{1-\alpha} R_t^\alpha - w_t L_t - p_{rt}(1 - s_{rt})R_t, \tag{11}$$

and the Lagrangian for energy-producing firm i in period t (where λ_c and λ_d are Lagrange multipliers),

$$\begin{aligned}\mathcal{L}_{it} = {} & p_{rit}\left[(A_{cit}X_{cit})^\varepsilon + (A_{dit}D_{it})^\varepsilon\right]^{1/\varepsilon} - X_{cit} - (p_d + \tau_{dt})D_{it} \\ & - w_{zt}[(1+\delta_z s_{zct}^2)(1 - s_{zct})Z_{cit} + (1+\delta_z s_{zdt}^2)(1 - s_{zdt})Z_{dit}] \\ & - \lambda_{cit}\left[A_{cit} - F_c(\mathbf{K}_{t-1}, Z_{cit})\right] - \lambda_{dit}\left[A_{dit} - F_d(\mathbf{K}_{t-1}, Z_{dit})\right],\end{aligned} \tag{12}$$

and take first-order conditions. These conditions—equations A.2–A.10, shown in Appendix A.1—yield the following equations in symmetric equilibrium, necessary conditions for an

[13] Why not assume that deadweight losses are simply in proportion to tax payments? The reason is that this assumption may lead to the perverse result that leakage leads to higher taxes rather than lower, since a higher tax causes energy producers to switch to clean inputs, and may therefore lead to lower tax payments even though the costs imposed on the producers (and hence the incentive to shift production abroad) are greater.

internal optimum:

$$\frac{X_{ct}}{(p_d+\tau_{dt})D_t} = \left(\frac{A_{ct}X_{ct}}{A_{dt}D_t}\right)^{\varepsilon}; \tag{FOC1}$$

$$\frac{\lambda_{ct}A_{ct}}{\lambda_{dt}A_{dt}} = \left(\frac{A_{ct}X_{ct}}{A_{dt}D_t}\right)^{\varepsilon}; \tag{FOC2}$$

$$(1+\delta_z s_{zct}^2)(1-s_{zct})w_{zt}Z_{ct} = \eta_{Aczt}\lambda_{ct}A_{ct} = \eta_{Aczt}X_{ct}; \tag{FOC3}$$

$$\frac{(1+\delta_z s_{zct}^2)}{(1+\delta_z s_{zdt}^2)}\frac{(1-s_{zct})}{(1-s_{zdt})}\frac{Z_{ct}}{Z_{dt}} = \frac{\eta_{Aczt}}{\eta_{Adzt}}\frac{\lambda_{ct}A_{ct}}{\lambda_{dt}A_{dt}}; \tag{FOC4}$$

$$\left(\frac{R_t}{A_{Lt}L}\right)^{1-\alpha} = \frac{\eta}{1-s_{rt}}\alpha\left[A_{ct}^{\varepsilon/(1-\varepsilon)} + \left(\frac{A_{dt}}{p_d+\tau_{dt}}\right)^{\varepsilon/(1-\varepsilon)}\right]^{(1-\varepsilon)/\varepsilon}. \tag{FOC5}$$

Here η_{Acz} is the elasticity of A_c with respect to investment Z_c, (and similarly for η_{Adz}) so

$$\eta_{Aczt} = \frac{\partial A_{ct}}{\partial Z_{ct}}\frac{Z_{ct}}{A_{ct}}.$$

The interpretation of these equations is as follows. FOC1 and FOC2 show that relative factor costs and the relative values of factor-augmenting knowledge stocks are equal, and (since ε is positive) the abundant factor accounts for the major share of costs. FOC3 shows that the marginal cost of research should be equal to its marginal benefit, and FOC4 shows that investments in factor-augmenting knowledge are in proportion to the shadow values of knowledge, modified by the relative elasticities of knowledge to investment (a sector with more elastic knowledge attracting higher investment). These results imply that—when the elasticities are equal—investments are in proportion to factor shares, in accordance with Hart (2013). Finally, FOC5 shows that effective aggregate energy inputs increase with effective labour inputs and the productivities of the energy inputs, and decrease in the emissions tax τ_d.

On the basis of the above equations we define a market equilibrium as follows.

Definition 1. *A market equilibrium of the model economy consists of a resource allocation $[X_{ct}, D_t, Z_{ct}, Z_{dt}]_{t=0}^{T}$ such that—given an initial state of knowledge $\mathbf{K}_{-1}$, initial pollution stock S_{-1}, and regulation $[s_{rt}, \tau_{dt}, s_{zct}, s_{zdt}]_{t=0}^{T}$—the relative quantities of energy inputs X_{ct} and D_t are in accordance with FOC1, total investment is in accordance with FOC3, relative investment is in accordance with FOC2 and FOC4, effective energy inputs R_t are in accordance with FOC5, and energy-augmenting knowledge stocks grow in accordance with equations (7)–(9).*

3.2. The social planner's solution and regulation in first best

Having characterized the equilibrium of the market economy we now turn to the social planner's solution, i.e. the allocation which would be chosen by a benevolent planner who could simply allocate resources as she wished, without deadweight losses. We use the FOCs on the planner's problem to draw conclusions about optimal regulation in first-best (Propositions 1 and 2). These provide a theoretical understanding of when and why subsidies to clean research may be larger than subsidies to fossil research in an optimally regulated economy in which the regulator has access to sufficient instruments—without deadweight losses—to achieve the first-best (planner's) allocation.

The benevolent social planner chooses an allocation to maximize U, subject to the following constraints: gross consumption X_t is equal to gross production Y_t minus extraction inputs X_{ct} and $p_d D_t$, and research costs $w_z(Z_c + Z_d)$; fossil use D_t leads to emissions P_t; emissions P_t add to the pollution stock S_t; and research effort boosts the productivities A_c and A_d. The Lagrangian of this maximization problem follows, where we do not include the fossil stock restriction because we assume it is not binding:

$$\begin{aligned}\mathcal{L} = \sum_{t=0}^{\infty}\beta^t \Big\{ & (A_{Lt}L)^{1-\alpha}\left[(A_{ct}X_{ct})^{\varepsilon} + (A_{dt}D_t)^{\varepsilon}\right]^{\alpha/\varepsilon} - X_{ct} - p_d D_t \\ & - w_{zt}(Z_{ct} + Z_{dt}) \\ & + \nu_t(P_t - D_t) + \xi_t\left[S_t - G(H_t) - P_t\right] \\ & - \mu_{ct}\left[A_{ct} - F_c(\mathbf{K}_{t-1}, Z_{ct})\right] - \mu_{dt}\left[A_{dt} - F_d(\mathbf{K}_{t-1}, Z_{dt})\right] \Big\} \exp(-\gamma_1 S_t^{\gamma_2}).\end{aligned} \tag{13}$$

There are four Lagrange multipliers: ν is the shadow price of polluting emissions, ξ is the shadow price of the pollution stock, and μ_c and μ_d are the shadow prices of the knowledge stocks.[14]

Take first-order conditions on the control variables—i.e. X_c, D, Z_c and Z_d—to yield the following necessary conditions for the planner (corresponding to FOCs 1, 3 and 4 and 5 above):

$$\frac{X_{ct}}{(p_d + \nu_{dt})D_t} = \left(\frac{A_{ct}X_{ct}}{A_{dt}D_t}\right)^{\varepsilon}; \tag{FOC1$'$}$$

$$w_{zt}Z_{ct} = \eta_{Aczt}\mu_{ct}A_{ct}; \tag{FOC3$'$}$$

$$\frac{Z_{ct}}{Z_{dt}} = \frac{\eta_{Aczt}}{\eta_{Adzt}}\frac{\mu_{ct}A_{ct}}{\mu_{dt}A_{dt}}; \tag{FOC4$'$}$$

$$\left(\frac{R_t}{A_{Lt}L}\right)^{1-\alpha} = \alpha\left[A_{ct}^{\varepsilon/(1-\varepsilon)} + \left(\frac{A_{dt}}{p_d + \tau_{dt}}\right)^{\varepsilon/(1-\varepsilon)}\right]^{(1-\varepsilon)/\varepsilon}. \tag{FOC5$'$}$$

Compare the two sets of equations to obtain the following proposition.

[14]Since pollution affects utility rather than production—and therefore does not affect market prices directly—we define the Lagrange multipliers of the planner in the corresponding way, which is why the term $\exp(-\gamma_1 S_t^{\gamma_2})$ multiplies the entire expression on the RHS. This is simply a matter of definition, analogous to the difference between current value and present value formulations of dynamic problems.

Proposition 1. *The following equations are necessary conditions for first-best optimal regulation.*

$$\tau_{dt} = \nu_t; \quad (14)$$

$$1 - s_{rt} = \eta; \quad (15)$$

$$1 - s_{zct} = \lambda_{ct}/\mu_{ct}; \quad (16)$$

$$\frac{1 - s_{zct}}{1 - s_{zdt}} = \frac{\lambda_{ct}/\lambda_{dt}}{\mu_{ct}/\mu_{dt}}. \quad (17)$$

In addition, $\delta_\tau \nu_t = 0$ *and* $\delta_z(s_{zct} + s_{zdt}) = 0$*.*

Proof. A necessary condition for first-best optimal regulation is that the market allocation under optimal regulation (Definition 1) is identical to the allocation on the planner's optimal path. Furthermore, there should be no deadweight losses. Equations 14–17 then follow directly from comparison of FOC1–FOC5 and FOC1′–FOC5′ when $\delta_z = 0$. When $\delta_z > 0$ then first-best can be achieved if the $\lambda_c/\mu_c = \lambda_d/\mu_d = 1$, and hence research subsidies are zero, so the necessary conditions above still hold. □

Equation (14) shows that we have a standard Pigovian tax in first-best. Equation (15) shows that the subsidy to the energy intermediate is strictly positive, reflecting the presence of market power in the energy sector, which leads individual energy producers to restrict their production. Equation (16) shows that the proportion of research costs paid by firms themselves should be equal to the ratio of the firms' valuation of productivity to the planner's valuation; so if the planner values an increase in firm productivity twice as highly as the firm, then the regulator should pay 50 percent of firm research costs in first-best. The most interesting result is equation (17), which shows that subsidies to clean research relative to subsidies to fossil are equal to the ratio of (i) the relative shadow prices of the productivity indices for the representative firm, to (ii) the relative shadow prices of the productivity indices for the regulator. So clean research is more heavily subsidized than fossil research when clean research is undersupplied in the market relative to fossil research.

3.3. The path of the Pigovian tax and subsidies to clean-energy research

When is clean research undersupplied relative to fossil? We return to this question after establishing Lemma 1 on the path of the Pigovian tax and the consequences for clean and fossil energy. Throughout this section we follow Golosov et al. (2014) and Acemoglu et al. (2016) by setting $\gamma_2 = 1$, implying that marginal damages are only a function of gross production X (recall equation 1). Allowing for $\gamma_2 > 1$ here (which we do in the numerical model) would complicate the analysis without adding important insights.

Lemma 1. *The Pigovian tax* τ_d *increases without bound in first-best. Hence there must exist a time* T *beyond which clean energy dominates (in the sense that* $A_{cT}X_{cT} > A_{dT}D_T$*) as long* *as the unit cost of delivering clean energy $A_{ct}X_{ct}$ is bounded above for all X_c and t. This holds even when fossil stocks are infinite.*

Proof. First we prove that τ_{dt} increases without bound, then we prove the second result. We know (equation 14) that $\tau_d = \nu$, i.e. the Pigovian tax is equal to the shadow price of polluting emissions. Take the FOC on (13) in P_t to show that

$$\nu_t = \xi_t + \beta \xi_{t+1} \frac{\partial G_{t+1}}{\partial P_t} + \beta^2 \xi_{t+2} \frac{\partial G_{t+2}}{\partial P_t} + \cdots,$$

and then substitute in the FOC in S_t—which states that $\gamma_1 X_t = \xi_t$ when $\gamma_2 = 1$—to show that

$$\nu_t = \tau_{dt} = \gamma_1 \left(X_t + \beta X_{t+1} \frac{\partial G_{t+1}}{\partial P_t} + \beta^2 X_{t+2} \frac{\partial G_{t+2}}{\partial P_t} + \cdots \right). \tag{18}$$

Since X increases without bound as long as τ_{dt} is bounded above it follows directly that τ_{dt} increases without bound.

Now to the second result. Since τ_{dt} increases without bound, the net price of fossil inputs increases without bound. Fossil efficiency A_d is bounded above (by physical laws), so if augmented clean inputs A_cX_c are not to be preferred in finite time, the unit cost of these inputs must grow without bound. □

The intuition behind Lemma 1 is straightforward. Marginal damages—and hence also the Pigovian tax—are in proportion to gross production X. As labour productivity increases without bound, X increases without bound, and hence the Pigovian tax increases without bound. If the costs of clean energy are bounded, there is certain to be a transition from fossil to clean in first best.

What are the consequences of the switch from fossil to clean energy for research subsidization? To focus on this question we begin by defining four elasticities which determine how existing knowledge spills over to boost future knowledge, and rule out asymmetries in these elasticities which would otherwise cause fossil and clean-energy research to be treated differently. We define $\eta_{Ac/c}$ as the elasticity of A_{ct+1} to A_{ct}, and $\eta_{Ad/c}$ as the elasticity of A_{dt+1} to A_{ct}, with analogous definitions of $\eta_{Ad/d}$ and $\eta_{Ac/d}$. Thus (for instance)

$$\eta_{Ac/ct} = \frac{\partial A_{ct+1}}{\partial A_{ct}} \frac{A_{ct}}{A_{ct+1}}.$$

We then have Proposition 2.

Proposition 2. *Assume an economy without deadweight losses in which $\eta_{Ac/d} = \eta_{Ad/c} = 0$ and $\eta_{Ac/c} = \eta_{Ad/d}$. Then (i) if the social planner's allocation is a balanced growth path then $s_{zc} = s_{zd}$ in an optimally regulated economy; and (ii) if the social planner's allocation* *involves a monotonic increase over time in the factor share of input C relative to input D then $s_{zc} > s_{zd}$ in an optimally regulated economy.*

Proof. See Appendix B. □

The first result is in the spirit of Hart (2008), who derives a related result when research may either augment labour or reduce abatement costs. The intuition behind it is as follows: on a balanced growth path on which emissions are optimally taxed, clean and fossil research are equally undersupplied, hence there is no reason to favour one over the other through larger subsidies to one than the other.

The second result shows that if there is a technology transition in which one sector (e.g. the fossil sector) diminishes in importance while another (clean) increases, then research into the increasingly important (clean) sector should be subsidized relative to the other. The reason is that clean knowledge today spills over to clean knowledge tomorrow (and similarly for fossil to fossil), and this contribution cannot be captured by the firm. During a technology transition from fossil to clean the value of these spillovers is proportionately larger for clean-augmenting knowledge than for fossil-augmenting knowledge, implying that clean knowledge, ceteris paribus, should be subsidized. Metaphorically, each researcher in the clean-energy sector will have many future researchers standing on their shoulders during a phase in which the share of clean energy is growing, and this justifies a higher subsidy. This follows Hart (2008) (Proposition 4) and especially Heggedal (2015).

4. Optimal regulation in second best

Regulatory instruments are in reality imperfect, and in the model we capture this by assuming that their use causes deadweight losses, as discussed in the introduction and specified in Section 2.3. In Sections 3.2 and 3.3 above we set the relevant parameters to zero and focused on the first-best regulatory solution, while in this section we account for the deadweight losses and hence analyse second-best regulatory solutions. In the presence of deadweight losses we can no longer derive useful analytical results from the first-order conditions, hence the focus of this section is to define the regulator's problem, parameterize the model, and run policy experiments.

4.1. The regulator's problem

The regulator's problem is to choose the path of the regulatory instruments in order to maximize discounted net production. Discounted net production is final-good production minus input costs, including the cost of the unproductive activity induced by the use of research subsidies and the deadweight loss associated with Pigovian taxation:

$$\max_{[s_{rt}, s_{zct}, s_{zdt}, \tau_{dt}]_{t=0}^{T}} U = \sum_{t=0}^{T} \beta^t \Big[(A_{Lt}L)^{1-\alpha} R_t^{\alpha} - X_{ct} - p_d D_t - \delta_\tau (c_{rt} - c_{rt}^*) R_t \\ - w_{zt} ((1 + \delta_z s_{zct}^2) Z_{ct} + (1 + \delta_z s_{zdt}^2) Z_{dt}) \Big] \exp(-\gamma_1 S_t^{\gamma_2}), \tag{19}$$

where A_{Lt} evolves exogenously and the paths of the variables R_t, X_{ct}, D_t, Z_{ct}, Z_{dt}, and S_t are determined by the choices of optimizing firms which take the regulatory instruments as given. We specify the research input as labour, which is supplied elastically at the wage paid to labour in the final-good production sector (so $w_{zt} = w_t$). As long as labour in energy research is a very small proportion of total labour (in our model it is much less than 1 percent) this simple approach is effectively equivalent to the more complex approach of assuming fixed total labour and arbitrage between the production and research sectors.

The link between the regulatory instruments and firm choices is described by equations A.1–A.10 (Appendix A.1), equation C.5 (Appendix C), equations 7–9 (Section 2.2), and the stock equations defined below (21–22). See Appendix C for equations as used in the numerical model, and their derivation from the above. Note also that when the limit on fossil stocks is binding—as it is in some of our policy experiments, and in laissez faire—we must extend the model by adding the scarcity rent to the fossil price faced by energy firms; see Appendix C.3.

4.2. Parameterization

We start by parameterizing the final-good sector, then we tackle the thorniest problem, which is to find values for ϕ and δ_z determining the strength of the stepping-on-toes effect and deadweight losses due to research subsidies respectively. We complete the parameterization with the stock–decay and damage models, and the energy sector.

The final-good sector

In the final-good sector we assume that the growth rate of labour productivity is 2 percent per year (giving an overall annual growth rate of close to 2 percent, in line with the Maddison data and following Acemoglu et al. (2012) and Nordhaus's DICE 2007 calibration), and set the energy share (i.e. α) to 0.05. Following Nordhaus, we set the interest rate to 5.5 percent per year, giving $\beta = 1.055^{-10}$.

Stepping on toes and deadweight losses from research subsidies

Here we calibrate the stepping-on-toes effect discussed by Jones and Williams (1998) and the deadweight losses from research subsidization which prevent the achievement of first-best. Our approach is to specify a simple endogenous growth model without an energy sector, and calibrate it to yield (i) 2 percent annual growth (22 percent per 10-year period), and (ii) 5 percent of resources devoted to research ($s = 0.05$) in a regulatory optimum in which 30 percent of total research funding comes from the state ($s_z = 0.3$).[15] We then take the parameter values for ϕ and δ_z derived in the model and assume that the same values apply in the energy sector. Furthermore, we derive the following equation for ζ in terms of L:[16]

$$\zeta = (\delta_K + \theta)/[(sL)/(1 + \delta_z s_z^2)]^{\phi}. \tag{20}$$

The model is described in detail in Appendix D.1, and the results are presented in Table 1.

Table 1: Parameters fixing the effectiveness of research subsidies

β	θ	δ_K	ϕ	δ_z
0.5854	0.22	0.01	0.1857	1.14

Given that, in balanced growth, research subsidies are 30 percent, the value of δ_z implies a relatively modest 9 percent loss of efficiency in the research sector due to the subsidy. The choice of $\delta_K = 0.01$ implies that the decay rate of knowledge per 10-year period is 1 percent; this slow rate of decay captures the idea that old knowledge will become irrelevant over time given overall technological progress in the economy.

The value of ϕ—the elasticity of aggregate new knowledge to aggregate research effort—is more striking, at just 0.19: why does production of knowledge not double when inputs double, yielding a unit elasticity? The reason is straightforward. As Romer (1994) discusses, knowledge is a non-rival good, and understanding the implications of this fact is key to understanding the growth process. One implication—completely standard in growth theory—is that a non-rival good can be used by any number of agents simultaneously. Another, however, is that the rate of production of knowledge does not necessarily increase when inputs are duplicated. Consider a researcher in a laboratory investigating some question, starting at time t_0, and assume that at time t_1 the researcher finds the answer. Now add an *identical* researcher in an *identical* laboratory performing *identical* research, also starting at t_0. Clearly if the researchers are truly identical then they will find the same answer simultaneously, and the rate of increase in knowledge will be unchanged. In reality we know that researchers are not identical, nor do they work on identical questions, and these differences will give rise to an increasing relationship between the number of researchers and the rate of technological progress. However, there is no reason *a priori* to assume that ϕ is anywhere close to 1.

[15]The growth rate comes from Maddison data, and the research data comes from the National Science Foundation.

[16]Recall that ζ is the parameter controlling the productivity of the research input (see equations 9 and 10). The equation comes from equation D.4, given that $\delta_A = \delta_K$.

The idea that many researchers compete to come up with the same innovation is standard in the patent-race literature (see for instance Reinganum, 1982), but has received less attention in the growth literature. Indeed, in most of the growth literature, the elasticity represented by ϕ is assumed (by default) to be equal to 1 (see for instance all the models discussed by Acemoglu, 2009, in Chapters 13 and 14). But several authors in the growth literature—notably Stokey (1995) and Jones and Williams (2000)—have discussed the tendency for duplication of research effort, the latter denoting it as the *stepping-on-toes* effect and setting the elasticity which we denote as ϕ equal to 0.5. Venturini (2012) performs a series of estimations of a generalized growth function using US data, in all of which the estimated value of ϕ is less than 0.05.

The stock–decay model and damages

The stock model is set up to match model predictions for a timescale of up to 500 years into the future.[17] Subject to this caveat we are able to build a very simple model which is consistent with long-run balanced growth with constant emissions. The stock–decay model is as follows, where S is the atmospheric stock and O may be loosely interpreted as the oceanic stock:

$$S_{t+1} = D_{t+1} + (1-\delta_1)S_t + \delta_2 O_t; \tag{21}$$

$$O_{t+1} = (1-\delta_2-\delta_3)O_t + \delta_1 S_t. \tag{22}$$

Note that S is defined as the excess stock of carbon in the atmosphere, measured in ppmv, i.e. it is the total stock in ppmv minus 280 (the stock prior to anthropogenic emissions). We parameterize the damage model to match the damage function of the Nordhaus/DICE model. Note that we could have used other models—such as those of Nordhaus (2008) or Golosov et al. (2014)—without significantly affecting the results.

The parameters and starting values for the stock–decay and damage models are shown in Table 2. Regarding the stock–decay model, there is a rather rapid exchange of carbon between the atmospheric stock and the intermediate sink, and a slow loss of carbon from the intermediate sink out of the system. Marginal damages are almost independent of the stock,

[17]It is well known that the dynamics of the atmospheric carbon stock are very poorly approximated by assuming a constant decay rate of additions to the stock (or first-order linear kinetics). As carefully explained by Archer et al. (2009), the reason for this is that the primary sink for atmospheric CO_2 is the oceans, but CO_2 in the oceans does not disappear; dissolved CO_2 is in a dynamic equilibrium with atmospheric CO_2, and as the stock of dissolved CO_2 increases so does the flow *back* from the oceans to the atmosphere. The result is that the decay rate in a linear equation falls both over time (given a single pulse of emissions) and as cumulative emissions increase (given a continuous flow of emissions). The natural way to capture this feature is to add additional stocks, for instance an atmospheric stock and an oceanic stock. We do this in a very simple way, with the aim of replicating the behaviour of the atmospheric stock as modelled by Archer et al. (2009). However, we limit our ambition to replication of the properties of one of their models, over a timescale of up to 500 years, for relatively modest total emissions. Our model is parameterized to yield results consistent with the GENIE-16 CS model for the case of a 1000 Pg pulse for the first 500 years.

hence there is no climate catastrophe, just steadily increasing costs.

Table 2: Parameters and starting values for the stock–decay and damage models.

Stock–decay		Damage
Parameters	Starting values	Nordhaus
$\delta_1 = 0.095$	$S = 118.6$	$\gamma_1 = 2.2 \times 10^{-5}$
$\delta_2 = 0.095$	$O = 44.5$	$\gamma_2 = 1.25$
$\delta_3 = 0.024$		

The energy sector

We now turn to the energy sector: σ_1 and σ_2; η, ε, and δ_τ; and A_{c0}, A_{d0}, $\bar{A}_c$, $\bar{A}_d$, K^*, p_d, and Q_0.

We start with DTC and hence the parameters σ_1 and σ_2 in equations 9 and 10 (the equations for knowledge accumulation). The empirical literature on DTC in the energy sector has tended to focus on establishing the existence of an effect of input prices on input-augmenting knowledge (for instance Newell et al., 1999, Popp, 2002), and furthermore establishing that new sector-specific knowledge (fossil, clean) tends to build on existing knowledge within the sector (for instance Noailly and Smeets, 2015, Aghion et al., 2016). However, there is little or no guidance in the literature on the exact form of the knowledge production function; for discussions of this issue see Acemoglu (2002) and Hart (2013). We therefore set σ_1 and σ_2 to values we judge to be reasonable—0.7 and 0.1 respectively—and test alternatives in sensitivity analysis. The choice of σ_1 implies that the effect of technological progress in the final-good sector on the productivity of research into energy-augmenting knowledge is large, in line with the intuition suggested by the discussion on wind technology in Section 2.2. Recall that the parameter σ_2 can be interpreted as the likelihood that a discovery in the clean sector is also usable in the fossil sector, and vice versa; the low value reflects the idea that the majority of research in the respective sectors will not be relevant to the other, consistent with the analysis of authors such as Noailly and Smeets (2015) and Aghion et al. (2016).

The parameter η determines the degree of market power enjoyed by energy producers (when $\eta = 1$ there is perfect competition) and hence the profits made by energy firms. We assume that barriers to entry are small, so long-run profits should be small. Based on this assumption we find the following expression for η:

$$\eta = (1+\theta)/[1+\theta+\phi(\delta_K+\theta)]. \tag{23}$$

For the derivation and discussion see Appendix A.2.

We set $\varepsilon = 0.75$, implying that the elasticity of substitution between clean and fossil sources in generating energy services (given by $1/(1-\varepsilon)$) is 4. This parameter is difficult to estimate, for various reasons including the difficulty of separating short- and long-run elasticity of substitution, and the difficulty of knowing what the elasticity will be during the crucial phase when the fossil and clean sectors are of approximately equal size. We choose a relatively high value to reflect the idea that the difficulty of substituting clean for fossil as the proportion of clean increases will diminish over time thanks to new technologies in (for instance) energy storage.[18] And we set $\delta_\tau = 0.1$, so deadweight losses on implementing and enforcing the emissions tax are 10 percent of the costs that the tax imposes on firms.

It remains to specify the starting values of A_c and A_d, the limiting values $\bar{A}_c$ and $\bar{A}_d$, the level of K^* (thus determining the starting values of K_c and K_d), the level of p_d, and the initial fossil stock Q_0. Our aim is to find reasonable values which are broadly consistent with the available evidence. This means that our starting point should be consistent with a history of low investment in clean technology compared to fossil, and that the limits on the productivities of fossil and clean technology should be realistic, implying that (at current fossil extraction costs) the limiting cost of clean energy is broadly similar to that of fossil energy. Furthermore, the starting values must be consistent with observed emissions in 'period zero' (2005–2014), which is the period prior to the first period in the model. This process is described in Appendix D.2. The parameterization results in a current clean-energy price 2.67 times higher than the fossil-energy price, while in the limit (when knowledge approaches infinity but scarcity is zero) the costs of producing energy from clean and fossil inputs are equal.

4.3. Policy experiments

We now have a fully calibrated model, which we use to perform two policy experiments in our baseline parameterization, before turning to sensitivity analysis.

1. Policy experiment 1 (PE1):
   - First-best (hypothetical);
   - Second-best (optimal);
   - Laissez-faire.
2. Policy experiment 2 (PE2):
   - Second-best where regulator can only use emissions pricing;
   - Second-best where regulator can only use research subsidies;

[18] There is a lack of agreement in the literature about this parameter: for instance, Acemoglu et al. (2012) use two values, a low value of 3 and a high value of 10, whereas Golosov et al. (2014) choose a benchmark value of just 0.95. The problem is that until recently the consensus of existing research was that the elasticity is less than 1 (see Stern, 2012), but common sense tells us that it must be higher. (If the elasticity is less than 1 then when the cost of clean energy falls its factor share also falls!) This situation is remedied by Papageorgiou et al. (2017), who estimate the elasticity of substitution between clean and dirty energy inputs of macroeconomic production functions and put the elasticity in electricity generation at 2, and outside this sector 3.

- Second-best where regulator is a climate-change denier.

The purpose of PE1 is to investigate optimal regulation in our baseline (second-best) scenario, and to compare the path of economy in this case with (a) the hypothetical case in which first-best regulation is possible, and (b) laissez-faire. In PE2 we investigate the effect of restrictions on the behaviour of the regulator, in order to shed light on the relative importance of emissions pricing *contra* research subsidies, and on the overall importance of climate policy compared to optimal policy pursued by a climate-change denier. The results are presented in Figures 1 and 2, and Table 3. The figures show the paths of key variables, whereas in Table 3 we show the net present value of utility in each scenario.

Policy Experiment 1

The results of PE1 are illustrated in Figure 1. The difference between first-best and second-best allocations is small (dashed and continuous lines), whereas the allocation in laissez-faire is radically different, with a much later transition to clean energy (Figure 1(b)) and hence a much higher maximum atmospheric carbon concentration (Figure 1(a)). In the case of laissez-faire, the transition is driven by the exhaustion of fossil fuels (and hence rising scarcity rent) rather than policy. In Table 3 we see that the differences in the NPV of utility between all the scenarios are modest; this is a consequence of the high rate of discount and the conservative (Nordhausian) damage function. In the third column we see that the difference in utility between the second-best and first-best allocations is significant: if we consider the utility gain on moving from laissez-faire to the second-best optimum as 100 percent, then an extra 34 percent could be gained in the hypothetical first-best scenario without deadweight losses.

Regarding the emissions tax in second-best, in Fig. 1(c) we see that it rises monotonically over time, as we know from Lemma 1 must happen in first-best when the tax is equal to the marginal external cost of emissions (MEC). However, the emissions tax in second-best is not exactly equal to MEC, due to the presence of deadweight losses. It is initially well below MEC (which we call sub-Pigovian), because initially—when the tax only makes up a small part of the cost of fossil inputs—the elasticity of emissions to the tax is low, while the tax causes significant deadweight losses. Over time, the tax rises, hence the elasticity of emissions to the tax rises, and the optimal second-best tax approaches MEC. Our main focus is on the time path of the emissions tax rather than its level at any one time. However, it is worth noting that the initial level of the tax in first-best adds 15 percent to the price of fossil fuel. After allowing for our simplification of only considering one fossil fuel, this is broadly in line with Nordhaus (2008).[19]

[19] See for instance Nordhaus (2008) Table 5.4. Nordhaus' optimal price is USD 12 per ton of CO_2 in 2015, corresponding to USD 20 per ton of coal and USD 4.3 per barrel of oil. Taking the coal price as USD 50 per ton,

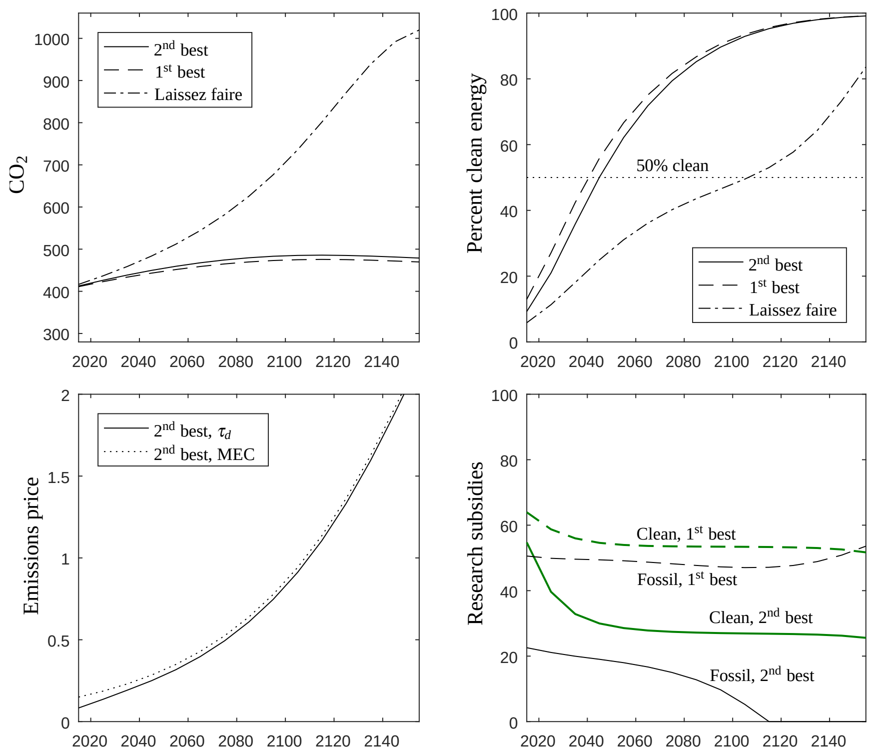

Figure 1: Transition paths in PE1: (a) Atmospheric carbon concentration, ppm; (b) Inputs of augmented clean energy as percentage of total, $A_cX_c/(A_cX_c + A_dD)$; (c) The emissions tax τ_d and the marginal external cost of emissions; (d) Percentage of research costs paid by subsidy in each sector (clean and fossil).

Turning to research subsidies, in Figure 1(d) we show subsidies for clean and fossil research effort, and compare them to the subsidies which would be applied in the absence of deadweight losses (dashed lines). In both cases clean research is subsidized more than fossil during the transition from fossil to clean technology, as expected from Proposition 2. The extra subsidization for clean research is greatest initially, when the clean 'pioneers' get relatively little benefit from their discoveries (because the clean share is small) but deliver big social benefits (because many future researchers will stand on their shoulders): the Heggedal effect. A notable feature is that the initial gap between clean and fossil subsidization is much greater in second-best than in the hypothetical first-best scenario. The reason is that in second-best the subsidy is raised partly because of the standing-on-shoulders effect described in Proposition 2, but also to compensate for the sub-Pigovian emissions taxes. We check this by running scenarios in which $\delta_\tau = 0$, so there are no deadweight losses from emissions taxes. Then the low initial tax disappears, as does part of the high initial subsidy.

Subsidies to fossil research fall gradually to zero in second-best, but are rather constant in first-best. The decline in second-best is because when fossil inputs are more productive, deadweight losses from fossil taxation are greater; more resources have to be put into en-

and the oil price as USD 50 per barrel, this implies that around 40 percent of the 2015 coal price is accounted for by the optimal tax, but just 8 percent of the oil price.

forcing compliance with the tax when the incentive to cheat is greater. So fossil efficiency —always a good thing in first-best—has a downside in second-best because it makes the clean transition harder to enforce.[20]

Table 3: Utilities (NPV over period 2015–2114) compared to laissez-faire.

	NPV utility (normalized)	Utility gain from policy (% of feasible max)
First-best, optimal policy	100.27	134
Second-best, optimal policy	100.20	100
Second-best, no research subsidies	100.18	91
Second-best, no emissions taxes	100.07	36
Second-best, climate-change denial	100.03	15
Laissez-faire within energy sector	100.00	0

Policy Experiment 2

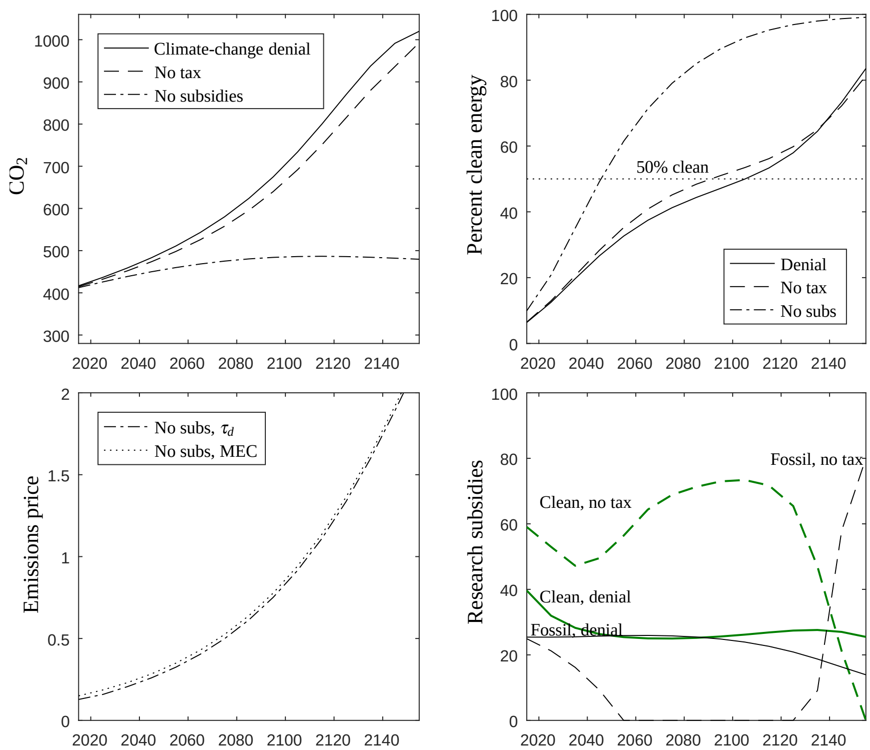

Figure 2: Transition paths in PE2: (a) Atmospheric carbon concentration, ppm; (b) Inputs of augmented clean energy as percentage of total, $A_cX_c/(A_cX_c + A_dD)$; (c) The emissions tax τ_d and the marginal external cost of emissions; (d) Percentage of research costs paid by subsidy in each sector (clean and fossil).

[20]Note also that subsidies to fossil research start to increase long after the clean transition. However, actual fossil research is extremely low at this time, and the reason for subsidizing it is that (a) it is then highly productive because there is little stepping-on-toes when there are few researchers, and (b) the ideas generated spill over to a small extent (with a one-period delay) to boost clean knowledge! So given our model of knowledge production and spillovers, subsidizing research in very small sectors may be advantageous because of high returns and spillovers to other (larger) sectors.

Recall that the purpose of PE2 is to shed light on the relative importance of emissions pricing and research subsidies, and on the overall importance of climate policy compared to optimal policy pursued by a climate-change denier. We start by looking at the effect of restricting the regulator by ruling out research subsidies, so the regulator must rely purely on emissions taxes. This restriction has only a small effect on the allocation, with the paths of atmospheric carbon and sectoral energy supply being similar to the second-best optimal. We can see this by comparing the 'dot–dash' curves in Figure 2 (a) and (b) with the corresponding second-best curves in Figure 1. Regarding emissions taxes, the degree to which the initial tax is sub-Pigovian is reduced, as we expect given the inability to subsidize research. The minor overall effect of the restriction is confirmed in Table 3, where we see that the restricted regulator (unable to use research subsidies) can achieve 91 percent of the utility increase which can be achieved by the unrestricted regulator.

Next we analyse the effect of ruling out emissions taxes, so the regulator must rely on research subsidies. This restriction has a very large effect on the allocation, with the paths of atmospheric carbon and sectoral energy supply being far from the optimal path, and instead close to laissez-faire (compare panels (a) and (b) in Figures 1 and 2, dashed lines in Figure 1 and dot–dash in Figure 2); as in laissez-faire, carbon stocks are exhausted and the atmospheric carbon concentration peaks at over 1000 ppm. The main difference between the laissez-faire and no-tax scenarios is that, in the latter, fossil consumption is somewhat delayed, a delay which is achieved through large and sustained subsidies to clean-energy research. On the other hand, subsidies to research into fossil-energy efficiency are actually lower in this scenario than in the unrestricted scenario of Figure 1. The reason is related to the reason for low fossil-research subsidies in PE1: fossil efficiency—always a good thing in first-best—has a major downside when the regulator is trying to encourage a transition to clean energy but is unable to tax fossil fuel inputs, because it makes fossil inputs more attractive.[21] Again, our conclusions about the effect of the restriction are borne out by the effect on utility shown in Table 3, where we see that the restricted regulator (unable to use emissions taxes) can achieve just 36 percent of the utility increase which can be achieved by the unrestricted regulator.

The third case in PE2 is when the regulator is assumed to be a climate-change denier, by which we mean that she attributes climate damages to forces beyond human control rather than carbon emissions. She therefore sets the shadow price of carbon emissions to zero. Now the path of atmospheric carbon and sectoral energy supply is extremely close to laissez-faire, and emissions taxes are (as expected) zero. Research subsidies are significant, and average subsidies to clean-energy research are higher than to fossil-energy research.

[21] Once the clean transition has happened and fossil scarcity has kicked in then the regulator is faced by a completely new situation with very scarce (expensive) fossil fuels and neglected fossil efficiency; this explains the dramatic rise in fossil-research subsidies after 2100.

Research is subsidized because of knowledge spillovers, and the subsidy to clean is greater than the subsidy to fossil during the transition from fossil to clean energy, which takes place throughout the period for which we present results in Figure 2. And in Table 3 we see that the research subsidization implemented by the climate-denying regulator yields an increase in utility which is just 15 percent of what can be achieved by an optimizing regulator.

4.4. Sensitivity analysis

The sensitivity analysis shows that the qualitative conclusions of the model are robust to the choice of parameters, and the effects of changes in parameters are for the most part predictable.

- When we double δ_z—doubling deadweight losses from research subsidization—research subsidies are greatly reduced (starting at 20 percent for clean-energy research and declining to 12 percent) while other results are broadly unchanged, as we expect given the relatively minor role played by research subsidies in the baseline case.
- When we double δ_τ—doubling deadweight losses from emissions taxation—the results illustrated in Figure 1(c) and (d) are accentuated: in short, the initial tax falls further, and the initial clean-research subsidy rises further. The timing of the clean transition is largely unaffected. On the other hand, as mentioned above, when we set the deadweight loss of emissions taxes to zero the low initial tax disappears, as does part of the high initial subsidy.
- When we reduce the long-run cost of clean energy by a factor of 1.25 there is less need for research subsidies, and an earlier transition.[22] The relative importance of the emissions tax (large) and research subsidies (smaller) remains broadly the same.
- When we reduce ε to 0.5 (reducing the elasticity of substitution from 4 to 2) the effect on the results is dramatic, because it means that significant fossil-fuel inputs are demanded even when the net price of fossil energy is considerably higher than clean energy. The result is that fossil-fuel stocks are gradually exhausted, and atmospheric carbon rises steadily beyond 600 ppm. Furthermore, because the transition from fossil to clean is much slower there is less need for the subsidy to clean-energy research, and clean and fossil-energy research subsidies track each other.
- When we reduce the spillover parameters, especially σ_1 which determines spillovers from labour-augmenting knowledge into the energy sector, this makes progress in clean-energy technology harder to achieve. The result is that the transition is delayed,

[22] When we increase the long-run cost of clean energy the reverse applies.

and the initial subsidy to clean research is increased, while there is little effect on the path of the emissions tax.

- If we increase the fossil stock, this affects the laissez-faire and no-tax scenarios in which the stock is exhausted. Utility in these scenarios goes down, because the scarcity rent (which acts as a proxy for the emissions tax) goes down.

5. Conclusions

In climate policy, the potential role of subsidies to clean research has recently gained increased attention (Acemoglu et al., 2012, 2016). In this paper we develop and calibrate a model to investigate the role of such subsidies: the basic structure of the model is intended to mirror that of IAMs such as Golosov et al. (2014), and it includes a model of directed technological change which is broadly consistent with historical data and physical restrictions regarding the energy sector and the overall economy.

We deliver analytical results in first best that clarify the roles of emissions prices and research subsidies. In Lemma 1 we show that given a standard damage function (as used by Golosov et al. (2014) and Acemoglu et al. (2016)) the Pigovian tax (and also the net price of using fossil inputs) increases without bound in a growing economy, because marginal damages are in proportion to gross production. This also implies that in first-best there is no change in the path of the optimal carbon tax when we compare models with DTC and with exogenously determined productivity. These results contrast strongly with those of Acemoglu et al. (2012) and Acemoglu et al. (2016). And in Proposition 2 we show that the motivation for higher subsidies to clean-energy research than to fossil-energy research in first best is that during a transition to clean energy each 'clean' researcher has more future researchers standing on her shoulders, an effect previously described by Heggedal (2015). These clarifications are important in the light of existing literature—such as Acemoglu et al. (2016)—in which the reader may gain the impression that a generally applicable motivation for high subsidies to clean research is that because such subsidies can redirect research effort, they should be used instead of a high carbon tax. This kind of motivation is only relevant in a second-best regulatory optimum when carbon pricing is either not available to the regulator, or are associated with deadweight losses.

We calibrate the model such that the starting point is broadly consistent with historical data, include deadweight losses caused by the use of policy instruments, and perform policy experiments and sensitivity analysis. The key conclusions are as follows. In the optimally regulated economy both emissions taxes and research subsidies are used, and there is a dramatic reduction in carbon emissions compared to laissez-faire. Subsidies to clean-energy research are initially large, and gradually decline as the transition to clean energy progresses. Simultaneously there are significant subsidies to fossil-energy research. The emissions tax starts low and increases monotonically, in line with what we expect from the analytical results; however, it is initially below the marginal external cost of emissions, due to deadweight losses and the low elasticity of emissions to the tax. The emissions tax is significantly more important than subsidies: when a regulator can only use the tax (and cannot subsidize research), 91 percent of the utility gain from optimal policy can still be achieved; when the regulator can only use research subsidies (and cannot tax emissions), only 36 percent of utility gains can be achieved. The sensitivity analysis shows that the qualitative results are robust to alternative parameter choices, and shifts in quantitative results are predictable.

The model includes several simplifications which are important for the results. We assume a single homogeneous fossil fuel, rather than following Golosov et al. (2014) and distinguishing between coal and oil/gas, or developing even more detailed models of input use. If we were to extend the model to allow for more inputs this could give further important insights, especially if we simultaneously extended the model to include multiple energy intermediates such as motive power (making things like cars move) and electricity supplied through the grid. We would then need at least six separate knowledge stocks: coal/grid and coal/motive power, oil/grid and oil/motive power, and finally clean/grid and clean/motive power. If we had empirical evidence about these stocks and their potential development, and about demand for the two intermediates, then we could find new and interesting results. The analysis of this paper suggests that we would probably find that a transition away from coal for electricity generation is already justified, implying that carbon taxes should be applied at close to the Pigovian level in this sector (i.e. equal to MEC), although not necessarily in other sectors where the transition occurs later.

We assume periods of 10 years, with technologies replaced each period. In reality many energy technologies (and the associated capital stocks) have much longer lifetimes than this. In an extended model with long-lived capital goods, firms' choices to invest in clean or fossil capital goods (such as power stations) would depend on their expectations about future policy, and if current policy were taken as a guide to the future this might mitigate against sub-Pigovian taxation in the present.

Appendix A. First-order conditions and market structure

Appendix A.1. First-order conditions

The first-order conditions in L and R_t on (11) yield

$$w_t = (1-\alpha)A_{Lt}^{1-\alpha}(R_t/L)^\alpha \tag{A.1}$$

and

$$p_{rt}(1-s_{rt}) = \alpha(A_{Lt}L/R_t)^{1-\alpha}. \tag{A.2}$$

Now turn to 12 and take FOCs and then assume symmetric equilibrium to yield the following equations. FOCs in inputs X_{ci} and D_i:

$$X_{ct} = \eta p_{rt} R_t^{1-\varepsilon} (A_{ct} X_{ct})^{\varepsilon} \tag{A.3}$$

and

$$(p_d + \tau_{dt}) D_t = \eta p_{rt} R_t^{1-\varepsilon} (A_{dt} D_t)^{\varepsilon}. \tag{A.4}$$

FOCs in Z_{ci} and Z_{di} (and using 7–9):

$$(1 + \delta_z s_{zct}^2)(1 - s_{zct}) w_{zt} Z_{ct} = \eta_{Aczt} \lambda_{ct} A_{ct} \tag{A.5}$$

and

$$(1 + \delta_z s_{zdt}^2)(1 - s_{zdt}) w_{zt} Z_{dt} = \eta_{Adzt} \lambda_{dt} A_{dt}, \tag{A.6}$$

where

$$\eta_{Aczt} = \frac{\partial A_{ct}}{\partial Z_{ct}} \frac{Z_{ct}}{A_{ct}} = \phi \frac{K_{ct} - (1 - \delta_K) K_{ct-1}}{(K^* + K_{ct}) K_{ct} / K^*} \tag{A.7}$$

and

$$\eta_{Adzt} = \frac{\partial A_{dt}}{\partial Z_{dt}} \frac{Z_{dt}}{A_{dt}} = \phi \frac{K_{dt} - (1 - \delta_K) K_{dt-1}}{(K^* + K_{dt}) K_{dt} / K^*}. \tag{A.8}$$

FOCs in A_{ci} and A_{di}:

$$\lambda_c A_{ct} = \eta p_{rt} R_t^{1-\varepsilon} (A_{ct} X_{ct})^{\varepsilon} \tag{A.9}$$

and

$$\lambda_d A_{dt} = \eta p_{rt} R_t^{1-\varepsilon} (A_{dt} D_t)^{\varepsilon}. \tag{A.10}$$

Obtaining equations FOC1–FOC5 is straightforward. FOC1 follows from (A.3) and (A.4), FOC2 follows from (A.9) and (A.10), FOC3 and FOC4 follow from (A.5)–(A.8). To find FOC5 is more involved. Rearrange (A.2), then find an expression for p_r in terms of A_c, A_d, p_d, and τ_d using (A.3), the production function for R (equation 3), and FOC1.

Appendix A.2. Market structure

Given the first-order conditions above, what are the profits made by each firm producing energy services? Firm revenue is $p_r R$, and input costs are the sum of X_c, $(p_d + \tau_d)D$, and research investments. Research investments change over time, thus causing profits to change. This should cause incentives for entry or exit, complicating the model. To deal with this problem, we first assume balance growth in which K_c and K_d are both much lower than K^*, implying (equation 6) that productivity A grows linearly in knowledge K in each sector, and with growth by a factor $1 + \theta$ per period. We can then insert $K_{ct-1} = K_{ct}/(1 + \theta)$ into A.7, and then (given that K_{ct}/K^* is small, and applying the same procedure in the other sector) we have

$$\eta_{Aczt} = \eta_{Aczt} = \phi(\theta + \delta_K)/(1 + \theta).$$

Now add together the input costs for the two energy inputs and research (using equations A.3–A.10) to yield

$$\text{Input costs } = p_r R\eta[1+\phi(\theta+\delta_K)/(1+\theta)].$$

Now we want to find the value of η which will ensure zero profits on this b.g.p. That is the value such that input costs are equal to revenue, $p_r R$. Hence we obtain equation 23.

This gives us a value for η that is consistent with free entry in the energy sector on a b.g.p. However, in general the economy will not be on a b.g.p.: as A_c and A_d approach their limits, research incentives will diminish, hence profits in the sector will rise and we would expect entry into the sector. An interesting extension would be to allow η to vary as a function of endogenous entry into the sector, approaching 1 when the mass of firms approached infinity. However, this is beyond our scope here, and instead we set η at the level consistent with balanced growth and low knowledge, and assume that it stays there due to barriers to entry.

Appendix B. Proof of Proposition 2

First use the FOC in A_{ct} for the market economy to show that

$$\lambda_{ct} A_{ct} = \alpha(A_{Lt}L)^{1-\alpha} R_t^{\alpha-\varepsilon}(A_{ct}C_t)^{\varepsilon},$$

and then show that the corresponding expression in the social planner's economy is

$$\mu_{ct} A_{ct} = \alpha(A_{Lt}L)^{1-\alpha} R_t^{\alpha-\varepsilon}(A_{ct}C_t)^{\varepsilon} + \beta\left[\mu_{ct+1}A_{ct+1}\eta_{Ac/c} + \mu_{dt+1}A_{dt+1}\eta_{Ad/c}\right].$$

Given that we have first-best regulation, the allocation in the planner's economy is the same as in the market economy, therefore we can simplify the above expression as follows:

$$\mu_{ct} A_{ct} = \lambda_{ct} A_{ct} + \beta\left[\mu_{ct+1}A_{ct+1}\eta_{Ac/ct} + \mu_{dt+1}A_{dt+1}\eta_{Ad/ct}\right]. \quad \text{(B.1)}$$

So the value of knowledge A_{ct} to the social planner is greater than its value to the representative firm, because knowledge in period t boosts knowledge in period $t+1$ (which in turn boosts knowledge in $t+2$, and so on).

(i) Assume that we are on a b.g.p. with a growth factor $1+\theta$ per period, and use (B.1) to show that

$$\mu_{ct} A_{ct} = \lambda_{ct} A_{ct} + \beta(1+\theta)\left[\mu_{ct}A_{ct}\eta_{Ac/ct} + \mu_{dt}A_{dt}\eta_{Ad/ct}\right].$$

Use the symmetric expression for $\mu_{dt}A_{dt}$ and the assumptions regarding the elasticities to show that $\lambda_c/\lambda_d = \mu_c/\mu_d$, and hence (using 17) that the optimal subsidy is zero.

(ii) Use (B.1) and the assumptions about elasticities to show that

$$\frac{\mu_{ct}A_{ct}}{\mu_{dt}A_{dt}} = \frac{\lambda_{ct}A_{ct} + \eta_{Ac/c}\beta\lambda_{ct+1}A_{ct+1} + (\eta_{Ac/c}\beta)^2\lambda_{ct+2}A_{ct+2} + \cdots}{\lambda_{dt}A_{dt} + \eta_{Ac/c}\beta\lambda_{dt+1}A_{dt+1} + (\eta_{Ac/c}\beta)^2\lambda_{dt+2}A_{dt+2} + \cdots}. \tag{B.2}$$

Given optimal regulation $\lambda_c A_c/(\lambda_d A_d)$ grows monotonically ((FOC1) and (FOC2)), so (B.2) implies that $\mu_{ct}A_{ct}/(\mu_{dt}A_{dt}) > \lambda_{ct}A_{ct}/(\lambda_{dt}A_{dt})$. But then (17) implies that $s_z > 0$, i.e. clean research should be subsidized.

Appendix C. Equations of the numerical model

Appendix C.1. The static model

We can solve the static model—in which we treat state variables as given—explicitly. The resulting equations are A.1 and A.2 (derived above), and

$$R = A_L L \left\{ \frac{\eta\alpha}{1-s_r} \left[A_c^{\varepsilon/(1-\varepsilon)} + \left(\frac{A_d}{p_d + \tau_d} \right)^{\varepsilon/(1-\varepsilon)} \right]^{(1-\varepsilon)/\varepsilon} \right\}^{1/(1-\alpha)} ; \tag{C.1}$$

$$X_c = R \left[A_c^\varepsilon + A_d^\varepsilon \left(\frac{A_c}{A_d} \right)^{-\varepsilon^2/(1-\varepsilon)} (p_d + \tau_d)^{-\varepsilon/(1-\varepsilon)} \right]^{-1/\varepsilon} ; \tag{C.2}$$

$$D = R \left[A_c^\varepsilon \left(\frac{A_c}{A_d} \right)^{\varepsilon^2/(1-\varepsilon)} (p_d + \tau_d)^{\varepsilon/(1-\varepsilon)} + A_d^\varepsilon \right]^{-1/\varepsilon} . \tag{C.3}$$

Equations C.1, C.2, and C.3 are derived using the definition of R (3) and equations A.1–A.4. First use A.3 and A.4 to yield

$$\frac{X_c}{D} = \left(\frac{A_c}{A_d} \right)^{\varepsilon/(1-\varepsilon)} (p_d + \tau_d)^{1/(1-\varepsilon)}. \tag{C.4}$$

Return to the definition of R to yield

$$R/D = [A_c^\varepsilon (X_c/D)^\varepsilon + A_d^\varepsilon]^{1/\varepsilon}.$$

Now take A.4, insert A.2 and rearrange to yield

$$D = R \left[\eta\alpha (A_L L/R)^{1-\alpha} A_d^\varepsilon \frac{1}{1-s_r} \frac{1}{p_d + \tau_d} \right]^{1/(1-\varepsilon)}$$

and hence

$$R = \left\{ \left(\frac{R}{D} \right)^{1-\varepsilon} \left[\eta\alpha (A_L L)^{1-\alpha} A_d^\varepsilon \frac{1}{1-s_r} \frac{1}{p_d + \tau_d} \right] \right\}^{1/(1-\alpha)} .$$

Insert the above expressions for R/D and X_C/D to yield

$$R = \left\{ \left[A_c^\varepsilon \left(\frac{A_c}{A_d} \right)^{\varepsilon^2/(1-\varepsilon)} (p_d + \tau_d)^{\varepsilon/(1-\varepsilon)} + A_d^\varepsilon \right]^{(1-\varepsilon)/\varepsilon} \eta \alpha (A_L L)^{1-\alpha} A_d^\varepsilon \frac{1}{1-s_r} \frac{1}{p_d + \tau_d} \right\}^{1/(1-\alpha)},$$

and hence

$$R = \left\{ \eta \alpha (A_L L)^{1-\alpha} A_d \frac{1}{1-s_r} \frac{1}{p_d + \tau_d} \left[\left(\frac{A_c}{A_d} \right)^{\varepsilon/(1-\varepsilon)} (p_d + \tau_d)^{\varepsilon/(1-\varepsilon)} + 1 \right]^{(1-\varepsilon)/\varepsilon} \right\}^{1/(1-\alpha)},$$

and finally rearrange to yield C.1. So we have an explicit solution for R in terms of state variables and the regulatory instruments s_r and τ_d. And using the equations for R/D and X_c/D above we can directly obtain equations C.2 and C.3.

Finally, we need the equations for c_{rt} and c_{rt}^*, unit costs of producing the energy intermediate with and without the tax. To find c_{rt}, divide total costs $X_c + p_d D$ by quantity R (from equation 3), then substitute for X_c/D using C.4. Rearrange to obtain

$$c_{rt} = \left\{ A_{ct}^{\varepsilon/(1-\varepsilon)} + [A_{dt}/(p_d + \tau_{dt})]^{\varepsilon/(1-\varepsilon)} \right\}^{-(1-\varepsilon)/\varepsilon}. \tag{C.5}$$

Then c_{rt}^* is found by setting $\tau_{dt} = 0$. (Note also that $c_{rt} = \eta p_{rt}$; unit costs are lower than the price because of market power.)

Appendix C.2. The dynamic model

Now we have the static solution, we must turn to the dynamic problem. Throughout these equations we substitute productivity A_c and A_d for knowledge K_c and K_d using equations 8, in order to reduce the number of variables in the numerical model.

The first two equations of the dynamic model follow directly from the production functions for knowledge, 9 and 10, after substituting for K_c and K_d using equations 8. The second two equations are as follows (note that we continue without explicit time subscripts, and use a subscript "$-$" to denote period $t-1$):

$$Z_c{}^{1-\phi} = \zeta \eta \phi L \frac{\alpha}{1-\alpha} \frac{A_{L-}^{\sigma_1} (K^*)^{1-\sigma_1} \left[A_{c-}/(\bar{A}_c - A_{c-}) + \sigma_2 A_{d-}/(\bar{A}_d - A_{d-}) \right]^{1-\sigma_1}}{(1+\delta_z s_{zc}^2)(1-s_{zc})(1-s_r)} \frac{(\bar{A}_c - A_c)^2}{\bar{A}_c A_c K^*} \frac{1}{1 + [(A_c/A_d)(p_d + \tau_d)]^{-\varepsilon/(1-\varepsilon)}}; \tag{C.6}$$

$$Z_d{}^{1-\phi} = \zeta \eta \phi L \frac{\alpha}{1-\alpha} \frac{(A_{L-}/K^*)^{\sigma_1} \left[A_{d-}/(\bar{A}_d - A_{d-}) + \sigma_2 A_{c-}/(\bar{A}_c - A_{c-}) \right]^{1-\sigma_1}}{(1+\delta_z s_{zd}^2)(1-s_{zd})(1-s_r)} \frac{(\bar{A}_d - A_d)^2}{\bar{A}_d A_d} \frac{1}{1 + [(A_c/A_d)(p_d + \tau_d)]^{\varepsilon/(1-\varepsilon)}}. \tag{C.7}$$

To derive these equations, take A.5, substitute for η_{Aczt} and λ_{ct}, then substitute for $wL/(p_r R)$ using equations A.1–A.2, and rearrange to yield

$$\frac{Z_c}{L} = \eta\phi \frac{1}{1+\delta_z s_{zc}^2} \frac{1}{1-s_{zc}} \frac{1}{1-s_r} \frac{\alpha}{1-\alpha} \frac{K_c - (1-\delta_K)K_{c-}}{(K^* + K_c)K_c/K^*} \left(\frac{A_c X_c}{R}\right)^{\varepsilon}.$$

Now substitute from the production function for knowledge and collect terms in Z_c to yield

$$Z_c{}^{1-\phi} = \eta\phi L \frac{1}{1+\delta_z s_{zc}^2} \frac{1}{1-s_{zc}} \frac{1}{1-s_r} \frac{\alpha}{1-\alpha} \frac{A_{L-}^{\sigma_1} (K_{c-} + \sigma_2 K_{d-})^{1-\sigma_1} \zeta}{(K^* + K_c)K_c/K^*} \left(\frac{A_c X_c}{R}\right)^{\varepsilon}.$$

Now substitute for K_c and K_d using 8, and for $A_c X_c/R$ using C.2, and then use symmetry, to yield C.6 and C.7.

Appendix C.3. Scarcity rent

When climate policy is sufficiently weak or non-existent the resource is exhausted asymptotically and there is a positive scarcity rent, which must increase by a factor $1/\beta$ per period (Hotelling). Equation 19 becomes

$$\max_{[s_{rt}, s_{zct}, s_{zdt}, \tau_{dt}]_{t=0}^T} U = \sum_{t=0}^{T} \beta^t \Big[(A_{Lt} L)^{1-\alpha} R_t^{\alpha} - X_{ct} - (p_d + \rho/\beta^t) D_t - \delta_\tau (c_{rt} - c_{rt}^*) R_t - w_{zt}((1+\delta_z s_{zct}^2) Z_{ct} + (1+\delta_z s_{zdt}^2) Z_{dt}) \Big] \exp(-\gamma_1 S_t^{\gamma_2}), \quad \text{(C.8)}$$

where ρ is the scarcity rent at $t=0$, the level of which is such that the resource is exhausted asymptotically. (This can be stated formally through a transversality condition.) Turning to the first-order conditions, the effect is that wherever we have $p_d + \tau_{dt}$ we replace it with $p_d + \rho/\beta^t + \tau_{dt}$.

Appendix D. Parameterization

Appendix D.1. Calibration of δ_z and ϕ

The calibration of δ_z and ϕ is based on the assumption that observed aggregate research subsidies s_z are optimally chosen by the regulator with perfect information. Furthermore, given these subsidies we observe a proportion s of aggregate resources devoted to research, and resultant growth by a factor $1+\theta$ per period. It builds on US data from the National Science Foundation[23] which shows that 3 percent of resources are devoted to research, and that 30 percent of total research funding comes from the state. We take the latter figure directly, whereas we revise the former up to 5 percent in our parameterization, for two reasons: firstly because some activities not recorded as research are likely to be research

[23] The data is at http://www.nsf.gov/statistics/natlpatterns/.

in the sense used in the model; and secondly because in the model we assume that the performance of research and the delivery of its results are simultaneous, whereas in reality research must be performed in advance. Assume that 3.6 percent of workers are engaged in research devoted to raising productivity in period t, and this research is performed in period $t-1$. Then the wage $w_{t-1} = w_t/(1+\theta)$, but the present value in period t of those wages is $w_t/[\beta(1+\theta)]$, and research costs are $0.05 w_t L$. We normalize $L = 1$, and follow our assumption from the main model by setting $\delta_A = 0.01$. And we assume (based on Maddison data) that long-run growth in per-capita GDP is 2 percent per year, so $\theta = 1.02^{10} - 1 = 0.22$.

Research		Growth		Labour	
s_z	0.3	θ	0.22	L	1
s	0.05	δ_A	0.01		

Table D.4: Input parameters to calibration model

The calibration model—without an energy sector—is as follows. Labour L is the only input to both production L_Y and research Z. Aggregate production is determined by effective labour,

$$Y_t = A_{Lt} L_Y,$$

while knowledge production in firm i is determined by effective research inputs Z_i:

$$A_{Lit} = (1-\delta_A)A_{Lt-1} + \zeta A_{Lt-1} Z_{it}^{\phi}. \tag{D.1}$$

As previously in the energy sector, there is a unit continuum of firms competing monopolistically to sell their differentiated products, and the demand elasticity for a given product is $1/(1-\eta)$. In the market economy firm i's problem is

$$\max_{L_{it}} \pi_{it} = A_{Lit}(Z_{lit})L_{it} - w_t L_{it} - w_t(1-s_{zt})(1+\delta_z s_{zt}^2)Z_{it},$$

subject to the knowledge production function, where s_z is the subsidy to research, and $\delta_z s_z^2 Z$ are labour inputs in the research sector which are not productive, and thus wasted. The Lagrangian for firm i in period t is as follows:

$$\mathcal{L}_{it} = A_{Lit}L_{it} - w_t L_{it} - w_t(1-s_{zt})(1+\delta_z s_{zt}^2)Z_{it} - \lambda_{it}\left[A_{Lit} - (1-\delta_A)A_{Lt-1} - \zeta A_{Lt-1} Z_{it}^{\phi}\right].$$

Take first-order conditions in L_{it}, Z_{lit}, and A_{Lit} and assume symmetric equilibrium to yield

$$Z_t^{1-\phi}(1-\delta_A+\zeta Z_t^{\phi}) = \frac{\phi\zeta L}{(1-s_z)(1+\delta_z s_z^2)}. \tag{D.2}$$

Assume balanced growth by a factor $1+\theta$ per period and use the knowledge production function (D.1) to show that

$$\zeta Z_t^{\phi} = \delta_A + \theta, \tag{D.3}$$

and given that the fraction of labour devoted to research is defined as s, so $(1+\delta_z s_z^2)Z/L = s$, we have the following expression for the research productivity parameter ζ:

$$\zeta = \frac{\delta_A + \theta}{[sL/(1+\delta_z s_z^2)]^{\phi}}. \tag{D.4}$$

Insert D.3 into D.2 and use the definition of s again to yield

$$\phi = s(1-s_z)\frac{1+\theta}{\delta_A+\theta} = 0.19. \tag{D.5}$$

Now that we know ϕ, equation D.4 gives us ζ as a function of δ_z. Which leaves us with the problem of finding δ_z.

To find δ_z we turn to optimal regulation. Our strategy is as follows: for any given δ_z, we know ζ and hence we know all the parameters of the model. Given these parameters we can find the optimal value of the subsidy s_z, which is unique as long as s_z is strictly decreasing in δ_z; we choose the unique value of δ_z for which $s_z = 0.3$. (And we know from the analysis of the firms' problem above that this will lead to the desired research allocation and growth rate.) The result is $\delta_z = 1.14$.[24]

Appendix D.2. Parameterization of starting values and limits in energy sector

We normalize the price of fossil fuels $p_d = 1$, and initial fossil productivity $A_{d0} = 1$. We then assume that 5 percent of augmented energy comes from renewable sources in period

[24]The regulator chooses s_z to boost Z. We therefore need the relationship between the two, which is given by equation D.2. We must then insert this into the regulator's utility function to find the optimal values of s_z and Z. To set up the utility function, first note that there is only one state variable—A_L—and that the regulator will always choose the same value of σ_z whatever the initial value of this variable. Hence the optimally regulated economy will always be on a balanced growth path (b.g.p.), and we can write the utility function as $U = \sum_{t=0}^{\infty} Y_0(Z)\beta^*(Z)^t$ where Y_0 is initial production, and β^* is the growth-adjusted discount factor:

$$Y_0(Z) = A_{L0}L[1 - Z(1+\delta_z s_z^2)/L];$$
$$\beta^*(Z) = (1 - \delta_A + \zeta Z^{\phi})\beta.$$

Furthermore, given balanced growth we can write the utility function as

$$U = Y_0(Z)(1-\beta^*(Z)).$$

All we need to do now to find the optimal s_z is to (numerically) take the locus of allowed points in (s_z, Z) space, and insert them into the utility function to plot utility as a function of s_z. If the value of s_z obtained is less than 0.3, this implies that our original choice of δ_z was too high, and vice versa. We adjust δ_z and iterate until we find the value of δ_z which gives an optimal $s_z = 0.3$.

zero. Then we can use FOC1 to find A_{c0} (which equals A_{c0}/A_{d0}):

$$A_{c0} = (5/95)^{(1-\varepsilon)/\varepsilon} = 0.3748.$$

So clean energy is initially 2.67 times more expensive than fossil.[25] Next we normalize $L = 1$, and use equations C.1–C.3 to find D_0 as a function of A_{L0},[26] and choose A_{L0} to deliver a period-zero rate of carbon emissions D_0 such that S (the excess stock, in ppmv) increases from 97.5 at the start of period zero (at the beginning of 2005) to 118.6 at the start of period one (end of 2014).[27] This implies (based on our stock–decay model) that $D_0 = 26.6$, which gives us $A_{L0} = 655$.

The next step is to find suitable initial knowledge stocks K_{c0} and K_{d0}. We require of these stocks that they are broadly consistent with historical energy shares. We start with K_{d0}, initial fossil-energy knowledge. Since the share of fossil energy in the overall economy has historically been around 5 percent, then (following Hart (2013)) around 5 percent of total research inputs should have gone to fossil-energy research historically. Given that we have A_{L0}, we can use equation 10 to find a value for K_{d0} which is consistent with such a rate of investment, assuming balanced growth and negligible spillovers of knowledge from clean to fossil:

$$K_{d0} = A_{L0}\left(\frac{\zeta\{\alpha Ls/[1+\delta_z s_z^2]\}^{\phi}}{\theta+\delta_K}\right)^{1/\sigma_1}.$$

To link K_{d0} to A_{d0} we need to know K^*. To find K^* we assume that fossil productivity A_{d0} is at 67 percent of the limit $\bar{A}_d$ (recall that the best coal-fired power stations are at 75 percent), so $A_{d0}/\bar{A}_d = 2/3$ (and $\bar{A}_d = 1.5$). This implies that $K_{d0}/(K^* + K_{d0}) = 2/3$, so

$$K^* = K_{d0}/2.$$

Our choice of $\bar{A}_c$ determines the lower limit on clean-energy costs relative to fossil costs (without allowing for scarcity rent or emissions taxes). We set the limiting costs equal, implying that $\bar{A}_c = \bar{A}_d$, since $p_d = 1$. The appropriate value here is very uncertain, depending crucially on the long-run prospects of cost reductions in renewable technologies.

[25]Note that this calibration involves a series of simplifications, inevitable given the simplified nature of the model. First, in reality there will not be a constant elasticity of substitution between renewable and fossil energy. Second, there are many types of non-fossil energy, including wind and solar but also nuclear; furthermore, there are also many types of fossil input (e.g. coal, oil, and gas). Third, the relative costs of augmented energy are not observable; even if we can observe spot prices for electricity produced from coal and wind, these prices are a result of the market interaction. Furthermore, energy intermediates have different qualities; in the model we aggregate electric power from the grid with motive power generated in a car's internal combustion engine.

[26]We know η (equation 23), and we assume that the period-zero tax on fossil fuels is zero, while the energy sector is subsidized at $s_r = 1 - \eta$.

[27]Thus we match the data from the Mauna Loa Observatory, Hawaii, given a baseline concentration of 280.

Our assumption here is optimistic compared to Acemoglu et al. (2016), who assume that long-run fossil costs are 85 percent of long-run renewable costs, assuming equal levels of technology.[28,29] And then we have $K_{c0} = K^* A_{c0}/(\bar{A}_c - A_{c0}) = 0.1665 K_{d0}$.

The initial stock of fossil fuels, Q_0, is irrelevant in the scenarios with optimal regulation, because fossil fuels are not exhausted. However, in scenarios when we restrict the behaviour of the regulator (for instance by assuming laissez-faire) then fossil fuels are exhausted, and we need a limit on the total quantity. We set this limit such that in our high-emissions scenarios with fossil exhaustion the atmospheric concentration peaks at around 1000 ppm. This is highly uncertain, and our choice is optimistic in the sense that fossil fuels are assumed to run out (or become too expensive to extract) at a lower level than predicted in the latest IPCC report.[30]

Acemoglu, Daron. 2002. Directed technical change. *Review of Economic Studies* 69 (4): 781–809.

__. 2009. *Introduction to modern economic growth*. Princeton University Press.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. The environment and directed technical change. *American Economic Review* 102 (1): 131–166.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. Transition to Clean Technology. *Journal of Political Economy* 124 (1): 52–104.

Aghion, Philippe and Peter Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60 (2): 323–51.

Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. 2016. Carbon taxes, Path Dependence and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy* 124 (1): 1–51.

Archer, David, Michael Eby, Victor Brovkin, Andy Ridgwell, Long Cao, Uwe Mikolajewicz, Ken Caldeira, Katsumi Matsumoto, Guy Munhoven, Alvaro Montenegro, and Kathy Tokos. 2009. Atmospheric Lifetime of Fossil Fuel Carbon Dioxide. *Annu. Rev. Earth Planet. Sci.* 37 (1): 117–134.

Arrow, Kenneth J. 1962. Economic Welfare and the Allocation of Resources for Invention, in R. Nelson, ed., *The Rate and Direction of Inventive Activity*, Princeton NJ: Princeton University Press.

[28] To derive the figure of 85 percent, consider the production functions for clean technology (unnumbered) and fossil (equation 2). Given that $\zeta = 0.016$ and $\nu = 0.04$ we have $\zeta^{\nu} = 0.85$, which is the ratio of fossil costs to clean costs at equal labour productivities.

[29] Note that our assumption of similar underlying cost structures in clean and fossil energy sectors implies that long-run shares of clean and fossil should be equal given zero taxes. However, if we assume that historical fossil prices were just 10 percent lower than future prices assumed in the model then this is consistent with a long-run renewable share of 1 percent in the absence of climate policy.

[30] See Figure 2.8, Pachauri and Meyer, eds (2014).

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

Fried, Stephie. 2018. Climate Policy and Innovation: A Quantitative Macroeconomic Analysis. *American Economic Journal: Macroeconomics* 10 (1): 90–118.

Gerlagh, Reyer. 2008. A climate-change policy induced shift from innovations in carbon-energy production to carbon-energy savings. *Energy Economics* 30 (2): 425–448.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica* 82 (1): 41–88.

Greaker, Mads, Tom-Reiel Heggedal, and Knut Einar Rosendahl. 2018. Environmental Policy and the Direction of Technical Change. *The Scandinavian Journal of Economics*, forthcoming.

Hart, Rob. 2008. The timing of taxes on CO_2 emissions when technological change is endogenous. *Journal of Environmental Economics and Management* 55 (2): 194–212.

_ . 2013. Directed technological change and factor shares. *Economics Letters* 119 (1): 77–80.

_ . 2016. Non-renewable resources in the long run. *Journal of Economic Dynamics and Control* 71: 1–20.

Hassler, John, Per Krusell, and Conny Olovsson. 2012. Energy-saving technical change. *NBER working paper series* (18456).

Heggedal, Tom-Reiel. 2015. Knowledge spillovers and R&D subsidies to new, emerging technologies. *Economics of Innovation and New Technology* 24 (7): 710–733.

Holman, Jack P. 1980. *Thermodynamics*, McGraw–Hill.

IEA. 2012. *Technology Roadmap: High-Efficiency, Low-Emissions Coal-Fired Power Generation*.

Jaimovich, Nir and Sergio Rebelo. 2017. Nonlinear Effects of Taxation on Growth. *Journal of Political Economy* 125 (1): 265–291.

Jones, Charles I. and John C. Williams. 1998. Measuring the social return to R&D. *Quarterly Journal of Economics* 113 (4): 1119–1135.

_ and _ . 2000. Too much of a good thing? The economics of investment in R&D. *Journal of Economic Growth* 5 (1): 65–85.

Klette, Tor Jakob, Jarle Møen, and Zvi Griliches. 2000. Do subsidies to commercial R&D reduce market failures? Microeconometric evaluation studies. *Research Policy* 29 (4–5): 471–495.

Lemoine, Derek. 2015. Energy Transitions: Directed Technical Change Meets Directed Extraction. Working Paper 15-04, University of Arizona.

Lipsey, R. G. and Kelvin Lancaster. 1956. The General Theory of Second Best. *The Review of Economic Studies* 24 (1): 11–32.

Maria, Corrado Di and Edwin van der Werf. 2008. Carbon leakage revisited: unilateral climate policy with directed technical change. *Environmental and Resource Economics* 39 (2): 55–74.

Newell, Richard G., Adam B. Jaffe, and Robert Stavins. 1999. The induced innovation hypothesis and energy-saving technological change. Quarterly Journal of Economics 114 (3): 941–975.

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

Nordhaus, William D. 1973. Some skeptical thoughts on theory of induced innovation. *Quarterly Journal of Economics* 87 (2): 208–219.

—. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.

Pachauri, Rajendra K. and Leo Meyer, eds. 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*. Geneva: IPCC.

Papageorgiou, Chris, Marianne Saam, and Patrick Schulte. 2017. Substitution between Clean and Dirty Energy Inputs: A Macroeconomic Perspective. *The Review of Economics and Statistics* 99 (2): 281–290.

Peretto, Pietro F. 2008. Effluent taxes, market structure, and the rate and direction of endogenous technological change. *Environmental and Resource Economics* 39 (2): 113–138.

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

Reinganum, Jennifer F. 1982. A Dynamic Game of R and D: Patent Protection and Competitive Behavior. *Econometrica* 50 (3): 671–688.

Romer, Paul M.. 1994. The origins of endogenous growth. *Journal of Economic Perspectives* 8 (1): 3–22.

Smulders, Sjak and Michiel de Nooij. 2003. The impact of energy conservation on technology and economic growth. *Resource and Energy Economics* 25 (1): 59–79.

Stern, David I. 2012. Interfuel Substitution: A Meta-analysis. *Journal of Economic Surveys* 26 (2): 307–331.

Stokey, Nancy L. 1995. R&D and economic growth. *Review of Economic Studies* 62 (3): 469–89.

Venturini, Francesco. 2012. Product variety, product quality, and evidence of endogenous growth. *Economics Letters* 117 (1): 74–77.