American Economic Review 2012, 102(1): 131–166
http://dx.doi=10.1257/aer.102.1.131

The Environment and Directed Technical Change†

By Daron Acemoglu, Philippe Aghion,
Leonardo Bursztyn, and David Hemous*

This paper introduces endogenous and directed technical change in a growth model with environmental constraints. The final good is produced from "dirty" and "clean" inputs. We show that: (i) when inputs are sufficiently substitutable, sustainable growth can be achieved with temporary taxes/subsidies that redirect innovation toward clean inputs; (ii) optimal policy involves both "carbon taxes" and research subsidies, avoiding excessive use of carbon taxes; (iii) delay in intervention is costly, as it later necessitates a longer transition phase with slow growth; and (iv) use of an exhaustible resource in dirty input production helps the switch to clean innovation under laissez-faire. (*JEL* O33, O44, Q30, Q54, Q56, Q58)

How to control and limit climate change caused by our growing consumption of fossil fuels and to develop alternative energy sources to these fossil fuels are among the most pressing policy challenges facing the world today.[1] While a large part of the discussion among climate scientists focuses on the effect of various policies on the development of alternative—and more "environmentally friendly"—energy sources, until recently the response of technological change to environmental policy has received relatively little attention by leading economic analyses of environment policy, which have mostly focused on computable general equilibrium models with exogenous technology.[2] Existing empirical evidence indicates that changes in the relative price of energy inputs have an important effect on the types of technologies that are developed and adopted. For example, Newell, Jaffe, and Stavins (1999)

* Acemoglu: Massachusetts Institute of Technology, Department of Economics, 50 Memorial Drive, Cambridge MA 02142-1347 and Canadian Institute for Advanced Research (e-mail: daron@mit.edu); Aghion: Harvard University, Department of Economics, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138; Stockholm University (IIES); and Canadian Institute for Advanced Research (e-mail: paghion@fas.harvard.edu); Bursztyn: University of California Los Angeles, Anderson School of Management, 110 Westwood Plaza, C-513 Los Angeles, CA 90095-1481 (e-mail: leonardo.bursztyn@anderson.ucla.edu); Hemous: Harvard University, Department of Economics, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138 (e-mail: hemous@fas.harvard.edu).
We thank Robert Barro, Emmanuel Farhi, Elhanan Helpman, Per Krusell, David Laibson, Ariel Pakes, Torsten Persson, Nicholas Stern, Nancy Stokey, Martin Weitzman, and three anonymous referees for very helpful suggestions. We also benefited from the comments of seminar and conference participants at Harvard, MIT, Stanford, Berkeley, IIES in Stockholm, Zurich, the NBER Summer Institute, the Midwest macro conference, the Canadian Institute for Advanced Research, the Latin American Meeting of the Econometric Society, TSE, and Simon Fraser University. Daron Acemoglu and Philippe Aghion gratefully acknowledge financial support, respectively, from the Toulouse Network for Information Technology (http://idei.fr/tnit/) and CIFAR, and from CIFAR and Bruegel.
† To view additional materials, visit the article page at http://dx.doi.org/10.1257/aer.102.1.131.

[1] See, for instance, Stott, Stone, and Allen (2004) on the contribution of human activity to the European heat wave of 2003, Emanuel (2005) and Landsea (2005) on the increased impact and destructiveness of tropical cyclones and Atlantic hurricanes over the last decades, and Nicholls and Lowe (2006) on sea-level rise.

[2] See, e.g., Nordhaus (1994), MacCracken et al. (1999), Nordhaus and Boyer (2000).

show that when energy prices were stable, innovations in air conditioning reduced the prices faced by consumers, but following the oil price hikes, air conditioners became more energy efficient. Popp (2002) provides more systematic evidence on the same point by using patent data from 1970 to 1994; he documents the impact of energy prices on patents for energy-saving innovations.

We propose a simple two-sector model of directed technical change to study the response of different types of technologies to environmental policies. A unique final good is produced by combining the inputs produced by these two sectors. One of them uses "dirty" machines and creates environmental degradation. Profit-maximizing researchers build on previous innovations ("build on the shoulders of giants") and direct their research to improving the quality of machines in one or the other sector.

Our model highlights the central roles played by the *market size* and the *price effects* on the direction of technical change (Acemoglu 1998, 2002). The market size effect encourages innovation towards the larger input sector, while the price effect directs innovation towards the sector with higher price. The relative magnitudes of these effects are, in turn, determined by three factors: (i) the elasticity of substitution between the two sectors; (ii) the relative levels of development of the technologies of the two sectors; (iii) whether dirty inputs are produced using an exhaustible resource. Because of the environmental externality, the decentralized equilibrium is not optimal. Moreover, the laissez-faire equilibrium leads to an "environmental disaster," where the quality of the environment falls below a critical threshold.

Our main results focus on the types of policies that can prevent such disasters, the structure of optimal environmental regulation and its long-run growth implications, and the costs of delay in implementing environmental regulation. Approaches based on exogenous technology lead to three different types of answers to (some of) these questions depending on their assumptions. Somewhat oversimplifying existing approaches and assigning colorful labels, we can summarize these as follows. The *Nordhaus* answer is that limited and gradual interventions are necessary. Optimal regulations should reduce long-run growth by only a modest amount. The Stern answer (see Stern 2009) is less optimistic. It calls for more extensive and immediate interventions, and argues that these interventions need to be in place permanently even though they may entail significant economic cost. The more pessimistic *Greenpeace* answer is that essentially all growth needs to come to an end in order to save the planet.

Our analysis suggests a different answer. In the empirically plausible case where the two sectors (clean and dirty inputs) are highly substitutable, immediate and decisive intervention is indeed necessary. Without intervention, the economy would rapidly head towards an environmental disaster, particularly because the market size effect and the initial productivity advantage of dirty inputs would direct innovation and production to that sector, contributing to environmental degradation. However, optimal environmental regulation, or even simple suboptimal policies just using carbon taxes or profit taxes/research subsidies, would be sufficient to *redirect* technical change and avoid an environmental disaster. Moreover, these policies need to be in place for only a *temporary* period, because once clean technologies are sufficiently advanced, research would be directed towards these technologies without further government intervention. Consequently, environmental goals can be achieved without permanent intervention and without sacrificing (much or any) long-run growth. While this conclusion is even more optimistic than Nordhaus's answer, as in the Stern or Greenpeace perspectives delay costs are significant, not simply because of the direct environmental damage, but because delay increases the technological gap between clean and dirty sectors, necessitating a more extended period of economic slowdown in the future.

Notably, our model also nests the Stern and Greenpeace answers. When the two sectors are substitutable but not sufficiently so, preventing an environmental disaster requires a permanent policy intervention. Finally, when the two sectors are complementary, the only way to stave off a disaster is to stop long-run growth.

A simple but important implication of our analysis is that optimal environmental regulation should always use both an input tax ("carbon tax") to control current emissions, and research subsidies or profit taxes to influence the direction of research. Even though a carbon tax would by itself discourage research in the dirty sector, using this tax both to reduce current emissions and to influence the path of research would lead to excessive distortions. Instead, optimal policy relies less on a carbon tax and instead involves direct encouragement to the development of clean technologies.

Our framework also illustrates the effects of exhaustibility of resources on the laissez-faire equilibrium and on the structure of optimal policy. An environmental disaster is less likely when the dirty sector uses an exhaustible resource (provided that the two sectors have a high degree of substitution) because the increase in the price of the resource as it is depleted reduces its use, and this encourages research towards clean technologies. Thus, an environmental disaster could be avoided without government intervention. Nevertheless, we also show that the structure of optimal environmental regulation looks broadly similar to the case without an exhaustible resource and again relies both on carbon taxes and research subsidies.

We illustrate some of our results with a simple quantitative example, which suggests that for high (but reasonable) elasticities of substitution between clean and dirty inputs (nonfossil and fossil fuels), the optimal policy involves an immediate switch of research and development to clean technologies. When clean and dirty inputs are sufficiently substitutable, the structure of optimal environmental policy appears broadly robust to different values of the discount rate (which is the main source of the different conclusions in the Stern report or in Nordhaus's research).

Our paper relates to the large and growing literature on growth, resources, and the environment. Nordhaus's (1994) pioneering study proposed a dynamic integrated model of climate change and the economy (the DICE model), which extends the neoclassical Ramsey model with equations representing emissions and climate change. Another branch of the literature focuses on the measurement of the costs of climate change, particularly stressing issues related to risk, uncertainty, and discounting.[3] Based on the assessment of discounting and related issues, this literature has prescribed either decisive and immediate governmental action (e.g., Stern 2007, in particular chapters 6–17) or a more gradualist approach (e.g., Nordhaus 2007), with modest control in the short-run followed by sharper emissions reduction in the medium and the long run. Recent work by Golosov et al. (2009) characterizes the

[3] For example, Stern (2007), Weitzman (2007, 2009), Dasgupta (2007, 2008), Nordhaus (2007), von Below and Persson (2008), Mendelsohn (2006), and Tol and Yohe (2006).

structure of optimal policies in a model with exogenous technology and exhaustible resources, where oil suppliers set prices to maximize discounted profits. They show that the optimal resource tax should be decreasing over time. Finally, some authors, for example, Hepburn (2006) and Pizer (2002), have built on Weitzman's (1974) analysis on the use of price or quantity instruments to study climate change policy and the choice between taxes and quotas.

Our paper is more closely related to the recent literature on the interactions between the environment, resources, and technology. Stokey (1998) shows that environment constraints can create an endogenous limit to growth, while Aghion and Howitt (1998, chapter 5) show that this may not be the case when "environment-friendly" innovations are allowed. Jones (2009) studies the conditions under which environmental and other costs of growth will outweigh its benefits. Early work by Bovenberg and Smulders (1995, 1996) and Goulder and Schneider (1999) study endogenous innovations in abatement technologies. Van der Zwaan et al. (2002) study the impact of environmental policies on technology in a model with learning-by-doing. Popp (2004) introduces directed innovation in the energy sector and presents a calibration exercise suggesting that models that ignore directed technical change might overstate the costs of environmental regulation. Gerlagh, Kverndokk, and Rosendahl (2009) also point out that using research subsidies would enable lower carbon taxes.[4] None of these works develops a systematic framework for the analysis of the impact of different types of environmental regulations on the direction of technical change.

The remainder of the article is organized as follows. Section I introduces our general framework. Section II focuses on the case without exhaustible resources. It shows that the laissez-faire equilibrium leads to an environmental disaster. It then shows how simple policy interventions can prevent environmental disasters and clarifies the role of directed technical change in these results. Section III characterizes the structure of optimal environmental policy in this setup. Section IV studies the economy with the exhaustible resource. Section V provides an example illustrating our results. Section VI concludes. Appendix I contains the proofs of some of the key results stated in the text, while Appendix II, which is available online, contains the remaining proofs and additional quantitative exercises.

I. General Framework

We consider an infinite-horizon discrete-time economy inhabited by a continuum of households comprising workers, entrepreneurs, and scientists. We assume that all households have preferences (or that the economy admits a representative household with preferences):

$$\sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} u(C_t, S_t), \tag{1}$$

[4] Other works investigating the response of technology to environment regulations include Grübler and Messner (1998), Goulder and Mathai (2000), Manne and Richels (2004), Messner (1997), Buonanno, Carraro, and Galeotti (2003), Nordhaus (2002), Sue Wing (2003), Di Maria and Valente (2006), Fischer and Newell (2008), Bosetti et al. (2008), Massetti, Carraro, and Nicita (2009), Grimaud and Rouge (2008), Gans (2009), and Aghion and Howitt (2009, chapter 16).

where C_t is consumption of the unique final good at time t, S_t denotes the *quality of the environment* at time t, and $\rho > 0$ is the discount rate.[5] We assume that $S_t \in [0, \overline{S}]$, where $\overline{S}$ is the quality of the environment absent any human pollution, and to simplify the notation, we also assume that this is the initial level of environmental quality, that is, $S_0 = \overline{S}$.

The instantaneous utility function $u(C,S)$ is increasing both in C and S, twice differentiable and jointly concave in (C,S). Moreover, we impose the following Inada-type conditions:

$$(2)\quad \lim_{C\downarrow 0} \frac{\partial u(C,S)}{\partial C} = \infty, \quad \lim_{S\downarrow 0} \frac{\partial u(C,S)}{\partial S} = \infty, \quad \text{and} \quad \lim_{S\downarrow 0} u(C,S) = -\infty.$$

The last two conditions imply that the quality of the environment reaching its lower bound has severe utility consequences. Finally we assume that

$$(3)\qquad \frac{\partial u(C,\overline{S})}{\partial S} = 0,$$

which implies that when S reaches $\overline{S}$, the value of the marginal increase in environmental quality is small. This assumption is adopted to simplify the characterization of optimal environmental policy in Section III.

There is a unique final good, produced competitively using "clean" and "dirty" inputs, Y_c and Y_d, according to the aggregate production function

$$(4)\qquad Y_t = \left(Y_{ct}^{(\varepsilon-1)/\varepsilon} + Y_{dt}^{(\varepsilon-1)/\varepsilon}\right)^{\varepsilon/(\varepsilon-1)},$$

where $\varepsilon \in (0, +\infty)$ is the elasticity of substitution between the two sectors and we suppress the distribution parameter for notational simplicity. Throughout, we say that the two sectors are (gross) *substitutes* when $\varepsilon > 1$ and (gross) *complements* when $\varepsilon < 1$ (throughout we ignore the Cobb-Douglas case of $\varepsilon = 1$).[6] The case of substitutes $\varepsilon > 1$ (in fact, an elasticity of substitution significantly greater than 1) appears as the more empirically relevant benchmark, since we would expect successful clean technologies to substitute for the functions of dirty technologies. For this reason, throughout the article we assume that $\varepsilon > 1$ unless specified otherwise (the corresponding results for the case of $\varepsilon < 1$ are discussed briefly in Section IID).

[5] For now, S can be thought of as a measure of general environmental quality. In our quantitative exercise in Section V, we explicitly relate S to the increase in temperature since preindustrial times and to carbon concentration in the atmosphere.

[6] The degree of substitution, which plays a central role in the model, has a clear empirical counterpart. For example, renewable energy, provided it can be stored and transported efficiently, would be highly substitutable with energy derived from fossil fuels. This reasoning would suggest a (very) high degree of substitution between dirty and clean inputs, since the same production services can be obtained from alternative energy with less pollution. In contrast, if the "clean alternative" were to reduce our consumption of energy permanently, for example by using less effective transport technologies, this would correspond to a low degree of substitution, since greater consumption of nonenergy commodities would increase the demand for energy.

The two inputs, Y_c and Y_d, are produced using labor and a continuum of sector-specific machines (intermediates), and the production of Y_d may also use a natural exhaustible resource:

$$Y_{ct} = L_{ct}^{1-\alpha}\int_0^1 A_{cit}^{1-\alpha}x_{cit}^{\alpha}\,di \text{ and } Y_{dt} = R_t^{\alpha_2}L_{dt}^{1-\alpha}\int_0^1 A_{dit}^{1-\alpha_1}x_{dit}^{\alpha_1}\,di, \tag{5}$$

where $\alpha, \alpha_1, \alpha_2 \in (0,1)$, $\alpha_1 + \alpha_2 = \alpha$, A_{jit} is the quality of machine of type i used in sector $j \in \{c,d\}$ at time t, x_{jit} is the quantity of this machine, and R_t is the flow consumption from an exhaustible resource at time t. The evolution of the exhaustible resource is given by the difference equation:

$$Q_{t+1} = Q_t - R_t, \tag{6}$$

where Q_t is the resource stock at date t. The per unit extraction cost for the exhaustible resource is $c(Q_t)$, where Q_t denotes the resource stock at date t, and c is a non-increasing function of Q. In Section IV, we study two alternative market structures for the exhaustible resource, one in which it is a "common resource" so that the user cost at time t is given by $c(Q_t)$, and one in which property rights to the exhaustible resource are vested with infinitely lived firms (or consumers), in which case the user cost will be determined by the Hotelling rule. Note that the special case where $\alpha_2 = 0$ (and thus $\alpha_1 = \alpha$) corresponds to an economy without the exhaustible resource, and we will first analyze this case.

Market clearing for labor requires labor demand to be less than total labor supply, which is normalized to 1, i.e.,

$$L_{ct} + L_{dt} \leq 1. \tag{7}$$

In line with the literature on endogenous technical change, machines (for both sectors) are supplied by monopolistically competitive firms. Regardless of the quality of machines and of the sector for which they are designed, producing one unit of any machine costs ψ units of the final good. Without loss of generality, we normalize $\psi \equiv \alpha^2$.

Market clearing for the final good implies that

$$C_t = Y_t - \psi\left(\int_0^1 x_{cit}\,di + \int_0^1 x_{dit}\,di\right) - c(Q_t)R_t. \tag{8}$$

The innovation possibilities frontier is as follows. At the beginning of every period, each scientist decides whether to direct her research to clean or dirty technology. She is then randomly allocated to at most one machine (without any congestion; so that each machine is also allocated to at most one scientist) and is successful in innovation with probability $\eta_j \in (0,1)$ in sector $j \in \{c,d\}$, where innovation increases the quality of a machine by a factor $1+\gamma$ (with $\gamma > 0$), that is, from A_{jit} to $(1+\gamma)A_{jit}$. A successful scientist, who has invented a better version of machine i in sector $j \in \{c,d\}$, obtains a one-period patent and becomes the entrepreneur for the current period in the production of machine i. In sectors where innovation is not successful, monopoly rights are allocated randomly to an entrepreneur drawn from the pool of potential entrepreneurs, who then uses the old technology.[7] This innovation possibilities frontier where scientists can target only a sector (rather than a specific machine) ensures that scientists are allocated across the different machines in a sector.[8] We also normalize the measure of scientists s to 1 and denote the mass of scientists working on machines in sector $j \in \{c, d\}$ at time t by s_{jt}. Market clearing for scientists then takes the form

$$s_{ct} + s_{dt} \leq 1. \tag{9}$$

Let us next define

$$A_{jt} \equiv \int_0^1 A_{jit}\, di \tag{10}$$

as the average productivity in sector $j \in \{c, d\}$, which implies that A_{dt} corresponds to "dirty technologies," while A_{ct} represents "clean technologies." The specification for the innovation possibilities frontier introduced above then implies that A_{jt} evolves over time according to the difference equation

$$A_{jt} = (1 + \gamma\eta_j s_{jt}) A_{jt-1}. \tag{11}$$

Finally, the quality of the environment, S_t, evolves according to the difference equation

$$S_{t+1} = -\xi Y_{dt} + (1 + \delta) S_t, \tag{12}$$

whenever the right-hand side of (12) is in the interval $(0, \overline{S})$. Whenever the right-hand side is negative, $S_{t+1} = 0$, and whenever the right-hand side is greater than $\overline{S}, S_{t+1} = \overline{S}$ (or, equivalently, $S_{t+1} = \min\{\max\langle -\xi Y_{dt} + (1 + \delta) S_t;\rangle\, 0; \overline{S}\}$). The parameter ξ measures the rate of environmental degradation resulting from the production of dirty inputs, and δ is the rate of "environmental regeneration." Recall also that $\overline{S}$ is the initial and the maximum level of environmental quality corresponding to zero

[7] The assumptions here are adopted to simplify the exposition and mimic the structure of equilibrium in continuous time models as in Acemoglu (2002) (see also Aghion and Howitt 2009 for this approach). We adopt a discrete time setup throughout to simplify the analysis of dynamics. Online Appendix II shows that the qualitative results are identical in an alternative formulation with patents and free entry (instead of monopoly rights being allocated to entrepreneurs).

[8] As highlighted further by equation (11) below, this structure implies that innovation builds on the existing level of quality of a machine and, thus, incorporates the "building on the shoulders of giants" feature. In terms of the framework in Acemoglu (2002, 2007), this implies that there is "state dependence" in the innovation possibilities frontier, in the sense that advances in one sector make future advances in that sector more profitable or more effective. This is a natural feature in the current context, since improvements in fossil fuel technology should not (and in practice do not) directly translate into innovations in alternative and renewable energy sources. Nevertheless, one could allow some spillovers between the two sectors, that is, "limited state dependence" as in Acemoglu (2002). In particular, in the current context, we could adopt a more general formulation which would replace the key equation (11) below by $A_{jt} = (1 + \gamma\eta_j s_{jt})\phi_j(A_{jt-1}, A_{\sim jt-1})$, for $j \in \{c, d\}$, where $\sim j$ denotes the other sector and ϕ_j is a linearly homogeneous function. Our qualitative results continue to hold provided that $\phi_c(A_c, A_d)$ has an elasticity of substitution greater than one as $A_c/A_d \rightarrow \infty$ (since in this case ϕ_c becomes effectively linear in A_c in the limit where innovation is directed at clean technologies).

pollution. This equation introduces the environmental externality, which is caused by the production of the dirty input.

Equation (12) encapsulates several important features of environmental change in practice. First, the exponential regeneration rate δ captures the idea that greater environmental degradation is typically presumed to lower the regeneration capacity of the globe. For example, part of the carbon in the atmosphere is absorbed by the ice cap; as the ice cap melts because of global warming, more carbon is released into the atmosphere, and the albedo of the planet is reduced, further contributing to global warming. Similarly, the depletion of forests reduces carbon absorption, also contributing to global warming. Second, the upper bound $\overline{S}$ captures the idea that environmental degradation results from pollution, and that pollution cannot be negative. We discuss below how our results change under alternative laws of motion for the quality of the environment.

Equation (12) also incorporates, in a simple way, the major concern of the majority of climate scientists, that the environment may deteriorate so much as to reach a "point of no return." In particular, if $S_t = 0$, then S_v will remain at 0 for all $v > t$. Our assumption that $\lim_{S \downarrow 0} u(C, S) = -\infty$ implies that $S_t = 0$ for any finite t cannot be part of a welfare-maximizing allocation (for any $\rho < \infty$). Motivated by this feature, we define the notion of an *environmental disaster*, which will be useful for developing the main intuitions of our model.

DEFINITION 1: *An environmental disaster occurs if* $S_t = 0$ *for some* $t < \infty$.

II. Environmental Disaster without Exhaustible Resources

In this and the next section, we focus on the case with $\alpha_2 = 0$ (and, thus, $\alpha_1 = \alpha$), where the production of the dirty input does not use the exhaustible resource. This case is of interest for several reasons. First, because the production technologies of clean and dirty inputs are symmetric in this case, the effects of directed technical change can be seen more transparently. Second, we believe that this case is of considerable empirical relevance, since the issue of exhaustibility appears secondary in several activities contributing to climate change, including deforestation and power generation using coal (where the exhaustibility constraint is unlikely to be binding for a long time). We return to the more general case where $\alpha_2 \neq 0$ in Section IV.

A. *The Laissez-Faire Equilibrium*

In this subsection we characterize the *laissez-faire equilibrium* outcome, that is, the decentralized equilibrium without any policy intervention. We first characterize the equilibrium production and labor decisions for given productivity parameters. We then analyze the direction of technical change.

DEFINITION 2: *An equilibrium is given by sequences of wages* (w_t), *prices for inputs* (p_{jt}), *prices for machines* (p_{jit}), *demands for machines* (x_{jit}), *demands for inputs* (Y_{jt}), *labor demands* (L_{jt}) *by input producers* $j \in \{c, d\}$, *research allocations* (s_{dt}, s_{ct}), *and quality of environment* (S_t) *such that, in each period t:* (i) (p_{jit}, x_{jit}) *maximizes profits by the producer of machine i in sector j;* (ii) L_{jt} *maximizes profits* *by producers of input j*; (*iii*) Y_{jt} *maximizes the profits of final good producers;* (*iv*) (s_{dt}, s_{ct}) *maximizes the expected profit of a researcher at date t*; (*v*) *the wage* w_t *and the prices* p_{jt} *clear the labor and input markets respectively; and* (*vi*) *the evolution of* S_t *is given by* (12).

To simplify the notation, we define $\varphi \equiv (1-\alpha)(1-\varepsilon)$ and impose the following assumption, which is adopted throughout the text (often without explicitly specifying it).

ASSUMPTION 1:

$$\frac{A_{c0}}{A_{d0}} < \min\left\{\left(1+\gamma\eta_c\right)^{-\frac{\varphi+1}{\varphi}}\left(\frac{\eta_c}{\eta_d}\right)^{\frac{1}{\varphi}}, \left(1+\gamma\eta_d\right)^{\frac{\varphi+1}{\varphi}}\left(\frac{\eta_c}{\eta_d}\right)^{\frac{1}{\varphi}}\right\}.$$

This assumption imposes the reasonable condition that initially the clean sector is sufficiently backward relative to the dirty (fossil fuel) sector that under laissez-faire the economy starts innovating in the dirty sector. This assumption enables us to focus on the more relevant part of the parameter space (Appendix I provides the general characterization).

We first consider the equilibrium at time t for given technology levels A_{cit} and A_{dit}. As the final good is produced competitively, the relative price of the two inputs satisfies

$$\frac{p_{ct}}{p_{dt}} = \left(\frac{Y_{ct}}{Y_{dt}}\right)^{-\frac{1}{\varepsilon}}. \tag{13}$$

This equation implies that the relative price of clean inputs (compared to dirty inputs) is decreasing in their relative supply and, moreover, that the elasticity of the relative price response is the inverse of the elasticity of substitution between the two inputs. We normalize the price of the final good at each date to one, i.e.,

$$\left[p_{ct}^{1-\varepsilon} + p_{dt}^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}} = 1. \tag{14}$$

To determine the evolution of average productivities in the two sectors, we need to characterize the profitability of research in these sectors, which will determine the direction of technical change. The equilibrium profits of machine producers endowed with technology A_{jit} can be written as (see Appendix I):

$$\pi_{jit} = (1-\alpha)\alpha p_{jt}^{1/(1-\alpha)} L_{jt} A_{jit}. \tag{15}$$

Taking into account the probability of success and using the definition of average productivity in (10), the expected profit Π_{jt} for a scientist engaged in research in sector j at time t is therefore

$$\Pi_{jt} = \eta_j(1+\gamma)(1-\alpha)\alpha p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1}. \tag{16}$$

Consequently, the relative benefit from undertaking research in sector *c* relative to sector *d* is governed by the ratio

$$(17)\qquad \frac{\Pi_{ct}}{\Pi_{dt}} = \frac{\eta_c}{\eta_d} \times \underbrace{\left(\frac{p_{ct}}{p_{dt}}\right)^{1/(1-\alpha)}}_{\text{price effect}} \times \underbrace{\frac{L_{ct}}{L_{dt}}}_{\text{market size effect}} \times \underbrace{\frac{A_{ct-1}}{A_{dt-1}}}_{\text{direct productivity effect}}.$$

The higher this ratio, the more profitable is R&D directed towards clean technologies. This equation shows that incentives to innovate in the clean versus the dirty sector machines are shaped by three forces: (i) the *direct productivity effect* (captured by the term A_{ct-1}/A_{dt-1}), which pushes towards innovating in the sector with higher productivity; this force results from the presence of the "building on the shoulders of giants" effect highlighted in (11); (ii) the *price effect* (captured by the term $\left(p_{ct}/p_{dt}\right)^{1/(1-\alpha)}$), encouraging innovation toward the sector with higher prices, which is naturally the relatively backward sector; (iii) the *market size effect* (captured by the term L_{ct}/L_{dt}), encouraging innovation in the sector with greater employment, and thus with the larger market for machines—when the two inputs are substitutes ($\varepsilon > 1$), this is also the sector with the higher aggregate productivity. Appendix I develops these effects more formally and shows that in equilibrium, equation (17) can be written as

$$(18)\qquad \frac{\Pi_{ct}}{\Pi_{dt}} = \frac{\eta_c}{\eta_d}\left(\frac{1 + \gamma\eta_c s_{ct}}{1 + \gamma\eta_d s_{dt}}\right)^{-\varphi-1}\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{-\varphi}.$$

The next lemma then directly follows from (18).

LEMMA 1: *Under laissez-faire, it is an equilibrium for innovation at time t to occur in the clean sector only when* $\eta_c A_{ct-1}^{-\varphi} > \eta_d(1 + \gamma\eta_c)^{\varphi+1}A_{dt-1}^{-\varphi}$, *in the dirty sector only when* $\eta_c(1 + \gamma\eta_d)^{\varphi+1}A_{ct-1}^{-\varphi} < \eta_d A_{dt-1}^{-\varphi}$, *and in both sectors when* $\eta_c(1 + \gamma\eta_d s_{dt})^{\varphi+1} A_{ct-1}^{-\varphi} = \eta_d(1 + \gamma\eta_c s_{ct})^{\varphi+1}A_{dt-1}^{-\varphi}$ (with $s_{ct} + s_{dt} = 1$).

PROOF:

See Appendix I.

The noteworthy conclusion of this lemma is that innovation will favor the more advanced sector when $\varepsilon > 1$ (which, in (18), corresponds to $\varphi \equiv (1 - \alpha)(1 - \varepsilon) < 0$).

Finally, output of the two inputs and the final good in the laissez-faire equilibrium can be written as

$$(19)\qquad Y_{ct} = (A_{ct}^{\varphi} + A_{dt}^{\varphi})^{-\frac{\alpha+\varphi}{\varphi}} A_{ct}A_{dt}^{\alpha+\varphi}, \qquad Y_{dt} = (A_{ct}^{\varphi} + A_{dt}^{\varphi})^{-\frac{\alpha+\varphi}{\varphi}} A_{ct}^{\alpha+\varphi}A_{dt},$$

$$\text{and } Y_t = (A_{ct}^{\varphi} + A_{dt}^{\varphi})^{-\frac{1}{\varphi}} A_{ct}A_{dt}.$$

Using these expressions and Lemma 1, we establish:

PROPOSITION 1: *Suppose that $\varepsilon > 1$ and Assumption 1 holds. Then there exists a unique laissez-faire equilibrium where innovation always occurs in the dirty sector only, and the long-run growth rate of dirty input production is $\gamma\eta_d$.*

PROOF:

See Appendix I.

Since the two inputs are substitutes ($\varepsilon > 1$), innovation starts in the dirty sector, which is more advanced initially (Assumption 1). This increases the gap between the dirty and the clean sectors and the initial pattern of equilibrium is reinforced: only A_d grows (at the rate $\gamma\eta_d > 0$), and A_c remains constant. Moreover, since φ is negative in this case, (19) implies that in the long run Y_d also grows at the rate $\gamma\eta_d$.

B. *Directed Technical Change and Environmental Disaster*

In this subsection, we show that the laissez-faire equilibrium leads to an environmental disaster and illustrate how a simple policy of "redirecting technical change" can avoid this outcome.

The result that the economy under laissez-faire will lead to an environmental disaster follows immediately from the facts that dirty input production Y_d always grows without bound (Proposition 1) and that a level of production of dirty input greater than $(1 + \delta)\xi^{-1}\overline{S}$ necessarily leads to a disaster next period. We thus have (proof omitted):

PROPOSITION 2: *Suppose that $\varepsilon > 1$ and Assumption 1 holds. Then the laissez-faire equilibrium always leads to an environmental disaster.*

Proposition 2 implies that some type of intervention is necessary to avoid a disaster. For a preliminary investigation of the implications of such intervention, suppose that the government can subsidize scientists to work in the clean sector, for example, using a proportional profit subsidy (financed through a lump-sum tax on the representative household).[9] Denoting this subsidy rate by q_t, the expected profit from undertaking research in the clean sector becomes

$$\Pi_{ct} = (1 + q_t)\eta_c(1 + \gamma)(1 - \alpha)\alpha p_{ct}^{\frac{1}{1-\alpha}} L_{ct} A_{ct-1},$$

while Π_{dt} is still given by (16). This immediately implies that a sufficiently high subsidy to clean research can redirect innovation towards the clean sector.[10] Moreover, while this subsidy is implemented, the ratio A_{ct}/A_{dt} grows at the rate $\gamma\eta_c$. When the two inputs are substitutes ($\varepsilon > 1$), a temporary subsidy (maintained for D periods)

[9] The results are identical with direct subsidies to the cost of clean research or with taxes on profits in the dirty sector.

[10] In particular, following the analysis in Appendix I, to implement a unique equilibrium where all scientists direct their research to the clean sector, the subsidy rate q_t must satisfy

$$q_t > (1 + \gamma\eta_d)^{-\varphi-1}\frac{\eta_d}{\eta_c}\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{\varphi} - 1 \text{ if } \varepsilon \geq \frac{2-\alpha}{1-\alpha}, \text{ and } q_t \geq (1 + \gamma\eta_c)^{(\varphi+1)}\frac{\eta_d}{\eta_c}\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{\varphi} - 1 \text{ if } \varepsilon < \frac{2-\alpha}{1-\alpha}.$$

is sufficient to redirect all research to the clean sector. More specifically, while the subsidy is being implemented, the ratio A_{ct}/A_{dt} will increase, and when it has become sufficiently high, it will be profitable for scientists to direct their research to the clean sector even without the subsidy.[11] Equation (19) then implies that Y_{dt} will grow asymptotically at the same rate as $A_{ct}^{\alpha+\varphi}$.

We say that the two inputs are *strong substitutes* if $\varepsilon \geq 1/(1-\alpha)$, or equivalently if $\alpha + \varphi \leq 0$. It follows from (19) that with strong substitutes, Y_{dt} will not grow in the long run. Therefore, provided that the initial environmental quality is sufficiently high, a temporary subsidy is sufficient to avoid an environmental disaster. This case thus delivers the most optimistic implications of our analysis: a temporary intervention is sufficient to redirect technical change and avoid an environmental disaster without preventing long-run growth or even creating long-run distortions. This contrasts with the Nordhaus, the Stern, and the Greenpeace answers discussed in the introduction.

If, instead, the two inputs are *weak substitutes*, that is $\varepsilon \in (1, 1/(1-\alpha))$ (or $\alpha + \varphi > 0$), then temporary intervention will not be sufficient to prevent an environmental disaster. Such an intervention can redirect all research to the clean sector, but equation (19) implies that even after this happens, Y_{dt} will grow at the rate $(1 + \gamma\eta_c)^{\alpha+\varphi} - 1 > 0$. Intuitively, since $\varepsilon > 1$, as the average quality of clean machines increases, workers are reallocated towards the clean sector (because of *the market size effect*). At the same time the increase of the relative price of the dirty input over time encourages production of the dirty input (*the price effect*). As shown in the previous paragraph, in the strong substitutes case the first effect dominates. In contrast, in the weak substitutes case, where $\varepsilon < 1/(1-\alpha)$, the second effect dominates,[12] and Y_{dt} increases even though A_{dt} is constant. In this case, we obtain the less optimistic conclusion that a temporary subsidy redirecting research to the clean sector will not be sufficient to avoid an environmental disaster; instead, similar to the Stern position, permanent government regulation is necessary to avoid environmental disaster. This discussion establishes the following proposition (proof in the text):

PROPOSITION 3: *When the two inputs are strong substitutes* $(\varepsilon \geq 1/(1-\alpha))$ *and* $\overline{S}$ *is sufficiently high, a temporary subsidy to clean research will prevent an environmental disaster. In contrast, when the two inputs are weak substitutes* $(1 < \varepsilon < 1/(1-\alpha))$, *a temporary subsidy to clean research cannot prevent an environmental disaster.*

This proposition shows the importance of directed technical change: temporary incentives are sufficient to *redirect* technical change towards clean technologies; with sufficient substitutability, once clean technologies are sufficiently advanced,

[11] The temporary tax needs to be imposed for D periods where D is the smallest integer such that
$$\frac{A_{ct+D-1}}{A_{dt+D-1}} > (1+\gamma\eta_d)^{\frac{\varphi+1}{\varphi}}\left(\frac{\eta_c}{\eta_d}\right)^{\frac{1}{\varphi}} \text{ if } \varepsilon \geq \frac{2-\alpha}{1-\alpha} \text{ and } \frac{A_{ct+D-1}}{A_{dt+D-1}} \geq (1+\gamma\eta_c)^{-\frac{\varphi+1}{\varphi}}\left(\frac{\eta_c}{\eta_d}\right)^{\frac{1}{\varphi}} \text{ if } 1 < \varepsilon < \frac{2-\alpha}{1-\alpha}.$$

[12] A different intuition for the $\varepsilon \in (1, 1/(1-\alpha))$ case is that improvements in the technology of the clean sector also correspond to improvements in the technology of the final good, which uses the clean intermediate as input; the final good, in turn, is an input for the dirty sector because machines employed in this sector are produced using the final good; hence, technical change in the clean sector creates a force towards the expansion of the dirty sector.

profit-maximizing innovation and production will automatically shift towards those technologies, and environmental disaster can be avoided without further intervention.

It is also useful to note that all of the main results in this section are a consequence of endogenous and directed technical change. Our framework would correspond to a model without directed technical change if we instead assumed that scientists are randomly allocated between the two sectors. Suppose, for simplicity, that this allocation is such that the qualities of clean and dirty machines grow at the same rate (i.e., at the rate $\gamma\tilde{\eta}$ where $\tilde{\eta} \equiv \eta_c \eta_d/(\eta_c + \eta_d)$). In this case, dirty input production will grow at the rate $\gamma\tilde{\eta}$ instead of the higher rate $\gamma\eta_d$ with directed technical change. This implies that when the two inputs are strong substitutes ($\varepsilon \geq 1/(1-\alpha)$), under laissez-faire a disaster will occur sooner with directed technical change than without. But while, as we have just seen, with directed technical change a temporary subsidy can redirect innovation towards the clean sector, without directed technical change such redirecting is not possible, and thus temporary interventions cannot prevent an environmental disaster.

C. *Costs of Delay*

Policy intervention is costly in our framework, partly because during the period of adjustment, as productivity in the clean sector catches up with that in the dirty sector, final output increases more slowly than the case where innovation continues to be directed towards the dirty sector. Before studying the welfare costs of intervention in detail in Section III, it is instructive to look at a simple measure of the (short-run) cost of intervention, defined as the number of periods T necessary for the economy under the policy intervention to reach the same level of output as it would have done within one period in the absence of the intervention: in other words, this is the length of the transition period or the number of periods of "slow growth" in output. This measure T_t (starting at time t) can be expressed as

$$T_t = \left\lceil \frac{\ln\left(\left((1+\gamma\eta_d)^{-\varphi} - 1\right)\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{\varphi} + 1\right)}{-\varphi \ln(1+\gamma\eta_c)} \right\rceil. \tag{20}$$

It can be verified that starting at any $t \geq 1$, we have $T_t \geq 2$ (in the equilibrium in Proposition 3 and with $\varepsilon \geq 1/(1-\alpha)$). Thus, once innovation is directed towards the clean sector, it will take more than one period for the economy to achieve the same output growth as it would have achieved in just one period in the laissez-faire equilibrium of Proposition 1 (with innovation still directed at the dirty sector). Then, the next corollary follows from equation (20) (proof omitted):

COROLLARY 1: *For $A_{dt-1}/A_{ct-1} \geq 1$, the short-run cost of intervention, T_t, is non-decreasing in the technology gap A_{dt-1}/A_{ct-1} and the elasticity of substitution ε. Moreover, T_t increases more with A_{dt-1}/A_{ct-1} when ε is greater.*

The (short-run) cost of intervention, T_t, is increasing in A_{dt-1}/A_{ct-1} because a larger gap between the initial quality of dirty and clean machines leads to a longer transition phase, and thus to a longer period of slow growth. In addition, T_t is also increasing in the elasticity of substitution ε. Intuitively, if the two inputs are close substitutes, final output production relies mostly on the more productive input, and therefore, productivity improvements in the clean sector (taking place during the transition phase) will have less impact on overall productivity until the clean technologies surpass the dirty ones.

The corollary shows that delaying intervention is costly, not only because of the continued environmental degradation that will result, but also because it will necessitate greater intervention; during the period of delay A_{dt}/A_{ct} will increase further, and thus when the intervention is eventually implemented, the duration of the subsidy to clean research and the period of slow growth will be longer. This result is clearly related to the "building on the shoulders of giants" feature of the innovation process. Furthermore, the result that the effects of ε and A_{dt-1}/A_{ct-1} on T are complementary implies that delaying the starting date of the intervention is more costly when the two inputs are more substitutable. These results imply that even though for the strong substitutes case the implications of our model are more optimistic than those of most existing analyses, immediate and strong interventions may still be called for.

Overall, the analysis in this subsection has established that a simple policy intervention that "redirects" technical change toward environment-friendly technologies can help prevent an environmental disaster. Our analysis also highlights that delaying intervention may be quite costly, not only because it further damages the environment (an effect already recognized in the climate science literature), but also because it widens the gap between dirty and clean technologies, thereby inducing a longer period of catch-up with slower growth.

D. *Complementary Inputs:* $\varepsilon < 1$

Although the case with $\varepsilon > 1$, in fact with $\varepsilon \geq 1/(1-\alpha)$, is empirically more relevant, it is useful to briefly contrast these with the case where the two inputs are complements, i.e., $\varepsilon < 1$. Lemma 1 already established that when $\varepsilon < 1$, innovation will favor the less advanced sector because $\varphi > 0$: in this case, the direct productivity effect is weaker than the combination of the price and market size effects (which now reinforce each other). Thus, under laissez-faire, starting from a situation where dirty technologies are initially more advanced than clean technologies, innovations will first occur in the clean sector until that sector catches up with the dirty sector; from then on innovation occurs in both sectors. Therefore, in the long run, the share of scientists devoted to the clean sector is equal to $s_c = \eta_d/(\eta_c + \eta_d)$, so that both A_{ct} and A_{dt} grow at the rate $\gamma\tilde{\eta}$. This implies that Proposition 2 continues to apply (see Appendix I).

It is also straightforward to see that a temporary research subsidy to clean innovation cannot avert an environmental disaster because it now has no impact on the long-run allocation of scientists between the two sectors, and thus A_{ct} and A_{dt} still grow at the rate $\gamma\tilde{\eta}$. In fact, $\varepsilon < 1$ implies that long-run growth is only possible if Y_{dt} also grows in the long run, which will in turn necessarily lead to an environmental disaster. Consequently, when the two inputs are complements ($\varepsilon < 1$), our model delivers the pessimistic conclusion, similar to the Greenpeace view, that environmental disaster can be avoided only if long-run growth is halted.

E. *Alternative Modeling Assumptions*

In this subsection, we briefly discuss the implications of a number of alternative modeling assumptions.

Direct Impact of Environmental Degradation on Productivity.—Previous studies have often used a formulation in which environmental degradation affects productivity rather than utility. But whether it affects productivity, utility, or both has little impact on our main results. Specifically, let us suppose that utility is independent of S_t, and instead, clean and dirty inputs $(j \in \{c, d\})$ are produced according to

$$Y_{jt} = \Omega(S_t) L_{jt}^{1-\alpha} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di, \tag{21}$$

where Ω is an increasing function of the environmental stock S_t, with $\Omega(0) = 0$. This formulation highlights that a reduction in environmental quality negatively affects the productivity of labor in both sectors. It is then straightforward to establish that in the laissez-faire equilibrium, either the productivity reduction induced by the environmental degradation resulting from the increase in A_{dt} occurs at a sufficiently high rate that aggregate output and consumption converge to zero, or this productivity reduction is not sufficiently rapid to offset the growth in A_{dt}, and an environmental disaster occurs in finite time. This result is stated in the next proposition (and proved in online Appendix II).

PROPOSITION 4: *In the laissez-faire equilibrium, the economy either reaches an environmental disaster in finite time or consumption converges to zero over time.*

With a similar logic to our baseline model, the implementation of a temporary subsidy to clean research in this case will avoid an environmental disaster and prevent consumption from converging to zero. It can also be shown that the short-run cost of intervention is now smaller than in our baseline model, since the increase in environmental quality resulting from the intervention also allows greater consumption.

Alternative Technologies.—First, it is straightforward to introduce innovations reducing the global pollution rate ξ or increasing the regeneration rate δ by various geoengineering methods. Since innovations in ξ or δ are pure public goods, there would be no research directed towards them in the laissez-faire equilibrium. This has motivated our focus on technologies that might be developed by the private sector.

Second, in our baseline model, dirty and clean technologies appear entirely separated. In practice, clean innovation may also reduce the environmental degradation resulting from (partially) dirty technologies. In fact, our model implicitly allows for this possibility. In particular, our model is equivalent to a formulation where there are no clean and dirty inputs, and instead, the unique final good is produced with the technology

$$Y_t = \left(\left(L_{ct}^{1-\alpha} \int_0^1 A_{cit}^{1-\alpha} x_{cit}^{\alpha} \, di \right)^{\frac{\varepsilon-1}{\varepsilon}} + \left(R_t^{\alpha_2} L_{dt}^{1-\alpha} \int_0^1 A_{dit}^{1-\alpha_1} x_{dit}^{\alpha_1} \, di \right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{22}$$

where A_{ct} and A_{dt} correspond to the fraction of "tasks" performed using clean versus dirty technologies, and the law of motion of the environmental stock takes the form

$$S_{t+1} = -\xi \times (Y_{dt}/Y_t) \times Y_t + (1 + \delta) S_t,$$

where Y_{dt}/Y_t measures the extent to which overall production uses dirty tasks. Clean innovation, increasing A_{ct}, then amounts to reducing the pollution intensity of the overall production process. Thus, our model captures one type of technical change that reduces pollution from existing production processes. We next discuss another variant of our model, also with pollution reducing innovations, which leads to similar but somewhat different results.

Substitution between Productivity Improvements and Green Technologies.—In our baseline model, clean technologies both increase output and reduce environmental degradation. An alternative is to remove the distinction between clean and dirty technologies and instead distinguish between technologies that increase the productivity of existing production methods and those that reduce pollution. Though this alternative reduces the ease with which the economy can switch to green technologies, many of the results are still similar.

To illustrate this point, suppose that the final good is produced according to the technology $Y_t = \int_0^1 A_{it}^{1-\alpha} x_{it}^{\alpha} di$ (i.e., in contrast to our baseline model, with only one type of machine), and the law of motion of the environment stock is given by $S_{t+1} = -\xi \int_0^1 e_{it}^{1-\alpha} x_{it}^{\alpha} di + (1 + \delta) S_t$, where e_{it} captures how dirty machine of type i is at time t. Research can now be directed either at increasing the productivity of machines, the A_{it}s, or at reducing pollution, the e_{it}s. Under laissez-faire, the equilibrium will again involve unbounded growth in output and an environmental disaster. However, an analysis similar to the one so far establishes that subsidies to innovations reducing pollution can redirect technical change and prevent such a disaster, though in this case such subsidies need to be permanent, and by reallocating research away from productivity improvements, they reduce long-run growth. The key reason why subsidies to clean research are less powerful in this case is that they are "complementary" to dirty technologies as reductions in e_{it} reduce the pollution from existing technologies instead of replacing them. We discuss the implications of this alternative technological assumption on the structure of optimal environmental regulation in Section IIIB (see online Appendix II for more details on these results).

Alternative Laws of Motion of Environmental Stock.—Several different variations of the laws of motion of the environmental stock, (12), yield similar results to our baseline model. For example, we could dispense with the upper bound on environmental quality, so that $\bar{S} = \infty$. In this case, the results are similar, except that a disaster can be avoided even if dirty input production grows at a positive rate, provided that this rate is lower than the regeneration rate of the environment, δ. An alternative is to suppose that $S_{t+1} = -\xi Y_{dt} + S_t + \Delta$, so that the regeneration of the environment is additive rather than proportional to current quality. With this alternative law of motion, it is straightforward to show that the results are essentially identical to the baseline formulation because a disaster can only be avoided if Y_{dt} does not grow at a positive exponential rate in the long run. Finally, the case in which pollution is created by the exhaustible resource will be discussed below.

III. Optimal Environmental Policy without Exhaustible Resources

We have so far studied the behavior of the laissez-faire equilibrium and discussed how environmental disaster may be avoided. In this section, we characterize the optimal allocation of resources in this economy and discuss how it can be decentralized using "carbon" taxes and research subsidies (we continue to focus on the case where dirty input production does not use the exhaustible resource, i.e., $\alpha_2 = 0$). The socially optimal allocation will "correct" for two externalities: (i) the environmental externality exerted by dirty input producers, and (ii) the knowledge externalities from R&D (the fact that in the laissez-faire equilibrium scientists do not internalize the effects of their research on productivity in the future). In addition, it will also correct for the standard static monopoly distortion in the price of machines, encouraging more intensive use of existing machines (see, for example, Aghion and Howitt 1998, or Acemoglu 2009). Throughout this section, we characterize a socially optimal allocation that can be achieved with lump-sum taxes and transfers (used for raising or redistributing revenues as required). A key conclusion of the analysis in this section is that optimal policy must use both a carbon tax (i.e., a tax on dirty input production) and a subsidy to clean research, the former to control carbon emissions and the latter to influence the path of future research. Relying on only carbon taxes would be excessively distortionary.

A. *The Socially Optimal Allocation*

The socially optimal allocation is a dynamic path of final good production Y_t, consumption C_t, input productions Y_{jt}, machine productions x_{jit}, labor allocations L_{jt}, scientist allocations s_{jt}, environmental quality S_t, and qualities of machines A_{jit} that maximizes the intertemporal utility of the representative consumer, (1), subject to (4), (5), (7), (8), (9), (11), and (12) (with $R_t \equiv 0$ and $\alpha_2 = 0$). The following proposition is one of our main results.

PROPOSITION 5: *The socially optimal allocation can be implemented using a tax on dirty input* (*a "carbon" tax*), *a subsidy to clean innovation*, *and a subsidy for the use of all machines* (*all proceeds from taxes/subsidies being redistributed/financed lump sum*).

PROOF:

See Appendix I.

This result is intuitive in view of the fact that the socially optimal allocation must correct for three market failures in the economy. First, the underutilization of machines due to monopoly pricing in the laissez-faire equilibrium is corrected by a subsidy for machines. Second, the environmental externality is corrected by introducing a wedge between the marginal product of dirty input in the production of the final good and its shadow value—which corresponds to a tax τ_t on the use of dirty input. In Appendix I (proof of Proposition 5), we show that

$$(23) \qquad \tau_t = \frac{\xi}{\hat{p}_{dt}} \frac{\frac{1}{1+\rho}\sum_{v=t+1}^{\infty}\left(\frac{1+\delta}{1+\rho}\right)^{v-(t+1)} I_{S_{t+1},\ldots,S_v<\bar{S}}\,\partial u(C_v,S_v)/\partial S}{\partial u(C_t,S_t)/\partial C},$$

where $\hat{p}_{jt}$ denotes the shadow (producer) price of input j at time t in terms of the final good (or more formally, as shown in Appendix I, it is the ratio of the Lagrange multipliers for constraints (5) and (4)), and $I_{S_{t+1},\ldots,S_v<\bar{S}}$ takes value 1 if $S_{t+1},\ldots,S_\nu<\bar{S}$ and 0 otherwise. This tax reflects that at the optimum, the marginal cost of reducing the production of dirty input by one unit must be equal to the resulting marginal benefit in terms of higher environmental quality in all subsequent periods. Finally, the socially optimal allocation also internalizes the knowledge externality in the innovation possibilities frontier and allocates scientists to the sector with the higher social gain from innovation. We show in Appendix I that in the social optimum, scientists are allocated to the clean sector whenever the ratio

$$(24) \qquad \frac{\eta_c(1+\gamma\eta_c s_{ct})^{-1}\sum_{v\geq t}\frac{\partial u(C_v,S_v)\,\partial C}{(1+\rho)^v}\hat{p}_{cv}^{1/(1-\alpha)}L_{cv}A_{cv}}{\eta_d(1+\gamma\eta_d s_{dt})^{-1}\sum_{v\geq t}\frac{\partial u(C_v,S_v)\,\partial C}{(1+\rho)^v}\hat{p}_{dv}^{1/(1-\alpha)}L_{dv}A_{dv}}$$

is greater than 1. This contrasts with the decentralized outcome where scientists are allocated according to the private value of innovation, that is, according to the ratio of the first term in the numerator over the first term in the denominator.[13]

That we need both a "carbon" tax and a subsidy to clean research to implement the social optimum (in addition to the subsidy to remove the monopoly distortions) is intuitive: the subsidy deals with future environmental externalities by directing innovation toward the clean sector, whereas the carbon tax deals more directly with the current environmental externality by reducing production of the dirty input. By reducing production in the dirty sector, the carbon tax also discourages innovation in that sector. However, using only the carbon tax to deal with both current environmental externalities and future (knowledge-based) externalities will typically necessitate a higher carbon tax, distorting current production and reducing current consumption excessively. An important implication of this result is that, without additional restrictions on policy, it is not optimal to rely only on a carbon tax to deal with global warming; one should also use additional instruments (R&D subsidies or a profit tax on the dirty sector) that direct innovation towards clean technologies, so that in the future production can be increased using more productive clean technologies.

[13] The knowledge externality is stark in our model because of the assumption that patents last for only one period. Nevertheless, our qualitative results do not depend on this assumption, since, even with perfectly enforced infinite-duration patents, clean innovations create a knowledge externality for future clean innovations because of the "building on the shoulders of giants" feature of the innovation possibilities frontier.

To elaborate on this issue, let us refer to optimal policy using both a carbon tax and a clean research subsidy as "first-best" policy, and to optimal policy constrained to use only the carbon tax as "second-best" policy (in both cases subsidies to the machines are present). Such a second-best policy might result, for example, because R&D subsidies are ineffective or their use cannot be properly monitored. Suppose first that both first-best and second-best policies result in all scientists being always allocated to the clean sector and that the first-best policy involves a positive clean research subsidy. In this case, we can show that the carbon tax in the second-best policy must be higher than in the first-best policy. This simply follows from the fact that under the second-best policy there is no direct subsidy to clean research, and thus the carbon tax needs to be raised to indirectly "subsidize" clean research. Nevertheless, when the clean research subsidy is no longer necessary in the first-best or in cases where under either the first-best or the second-best policies there is delay in the switch to clean research, carbon taxes may be lower for some time under the second-best policy than under the first-best policy (for example, because the switch to clean research may start later or finish earlier under the second-best).

B. *The Structure of Optimal Environmental Regulation*

In Section IIB, we showed that a switch to innovation in clean technologies induced by a temporary subsidy to clean research could prevent a disaster when the two inputs are substitutes. Here we show that, when the two inputs are sufficiently substitutable and the discount rate is sufficiently low, the optimal policy in Proposition 5 also involves a switch to clean innovation and only temporary taxes/subsidies (except for the subsidy correcting for monopoly distortions).

PROPOSITION 6: *Suppose that $\varepsilon > 1$ and the discount rate ρ is sufficiently small. Then all innovation switches to the clean sector in finite time, the economy grows asymptotically at the rate $\gamma\eta_c$, and the optimal subsidy on profits in the clean sector, q_t, is temporary. Moreover, if $\varepsilon > 1/(1-\alpha)$ (but not if $1 < \varepsilon < 1/(1-\alpha)$), then the optimal carbon tax, τ_t, is temporary.*

PROOF:

See online Appendix II.

To obtain an intuition for this proposition, first note that an optimal policy requires avoiding a disaster, since a disaster leads to $\lim_{S\downarrow 0} u(C,S) = -\infty$. This in turn implies that the production of dirty input must always remain below a fixed upper bound. When the discount rate is sufficiently low, it is optimal to have positive long-run growth, which can be achieved by technical change in the production of the clean input, without growth in the production of the dirty input (because $\varepsilon > 1$). Failing to allocate all research to clean innovation in finite time would then slow down the increase in clean input production and reduce intertemporal welfare. An appropriately chosen subsidy to clean research ensures that innovation occurs only in the clean sector, and when A_{ct} exceeds A_{dt} by a sufficient amount, innovation in the clean sector will have become sufficiently profitable that it will continue even after the subsidy is removed (and, hence, there is no longer a need for the subsidy).

The economy will then generate a long-run growth rate equal to the growth rate of A_{ct}, namely $\gamma\eta_c$. When $\varepsilon > 1/(1-\alpha)$, the production of the dirty input also decreases to 0 over time, and as a result, the environmental stock S_t reaches $\overline{S}$ in finite time due to positive regeneration. This in turn ensures that the optimal carbon tax given by (23) will reach zero in finite time.[14]

It is also straightforward to compare the structure of optimal policy in this model to the variant without directed technical change discussed briefly above. Since without directed technical change the allocation of scientists is insensitive to policy, redirecting innovation toward the clean sector is not possible. Consequently, optimal environmental regulation must prevent an environmental disaster by imposing an ever-increasing sequence of carbon taxes. This comparison highlights that the relatively optimistic conclusion that optimal environmental regulation can be achieved using temporary taxes/subsidies, and with little cost in terms of long-run distortions and growth, is a consequence of the presence of directed technical change.

Finally, it is also useful to return to the alternative modeling assumptions discussed in Section IIE, in particular, to the case where innovations can either increase the productivity of existing machines or reduce pollution. As already noted there, in this case, because clean technologies cannot directly replace dirty ones, subsidies to clean research need to be permanent. However, importantly, it can be shown that such subsidies to clean research, in addition to the standard carbon taxes, are again part of optimal environmental regulation even under this alternative technology, provided that either patents have finite (expected) duration or innovation creates knowledge spillovers (e.g., it involves creative destruction building on the shoulders of giants in the same variety as in our baseline model or it generates spillovers to other varieties; see online Appendix II for details).

IV. Equilibrium and Optimal Policy with Exhaustible Resources

In this section we characterize the equilibrium and the optimal environmental policy when dirty input production uses the exhaustible resource (i.e., when $\alpha_2 > 0$). In particular, we will show that the presence of an exhaustible resource may help prevent an environmental disaster because it increases the cost of using the dirty input even without policy intervention. Nevertheless, the major qualitative features of optimal environmental policy are similar to the case without exhaustible resource.

In the first two subsections, we simplify the exposition by assuming that there are no privately held property rights to the exhaustible resource. In this case, the user cost of the exhaustible resource is determined by the cost of extraction and does not reflect its scarcity value. We then show that the main results generalize to the case in which the property rights to the exhaustible resource are vested in infinite-lived firms or consumers, so that the price is determined by the Hotelling rule.

[14] This result depends on the assumption that $\partial u(C,\overline{S})/\partial S = 0$. With $\partial u(C,\overline{S})/\partial S > 0$, the optimal carbon tax may remain positive in the long run. Moreover, even under our assumptions, though temporary, optimal taxes/subsidies may sometimes be relatively long-lived, for example, as illustrated by our quantitative results in Section V. Finally, in practice the decline in carbon levels in the atmosphere is slower than implied by our simple equation (environment dynamics), necessitating a longer-lived carbon tax.

A. *The Laissez-Faire Equilibrium*

When $\alpha_2 > 0$, the structure of equilibrium remains mostly unchanged. In particular, the relative profitability of innovation in clean and dirty sectors reflects the same three effects as before: the direct productivity effect, the price effect, and the market size effect identified above. The only change relative to the baseline model is that the resource stock now affects the magnitude of the price and market size effects. In particular, as the resource stock declines, the effective productivity of the dirty input also declines, and its price increases, and the share of labor allocated to the dirty sector decreases with the extraction cost. The ratio of expected profits from research in the two sectors, which again determines the direction of equilibrium research, now becomes (see online Appendix II)

$$\frac{\Pi_{ct}}{\Pi_{dt}} = \kappa \frac{\eta_c c(Q_t)^{\alpha_2(\varepsilon-1)}}{\eta_d} \frac{(1 + \gamma\eta_c s_{ct})^{-\varphi-1}}{(1 + \gamma\eta_d s_{dt})^{-\varphi_1-1}} \frac{A_{ct-1}^{-\varphi}}{A_{dt-1}^{-\varphi_1}}, \tag{25}$$

where $\kappa \equiv \dfrac{(1 - \alpha)\alpha}{(1 - \alpha_1)\alpha_1^{(1+\alpha_2-\alpha_1)/(1-\alpha_1)}} \left(\dfrac{\alpha^{2\alpha}}{\psi^{\alpha_2}\alpha_1^{2\alpha_1}\alpha_2^{\alpha_2}}\right)^{(\varepsilon-1)}$

and $\varphi_1 \equiv (1 - \alpha_1)(1 - \varepsilon)$.

The main difference from the corresponding expression (18) in the case with $\alpha_2 = 0$ is the term $c(Q_t)^{\alpha_2(\varepsilon-1)}$ in (25). This new term, together with the assumption that $c(Q_t)$ is decreasing in Q_t, immediately implies that when the two inputs are substitutes ($\varepsilon > 1$), as the resource stock gets depleted, the incentives to direct innovations towards the clean sector will increase. Intuitively, the depletion of the resource stock increases the relative cost (price) of the dirty input and, thus, reduces the market for the dirty input and encourages innovation in the clean sector (because $\varepsilon > 1$). In fact, it is straightforward to see that asymptotically there will be innovation in the clean sector only (either because the extraction cost increases sufficiently rapidly, inducing all innovation to be directed at clean machines, or because the resource stock gets fully depleted in finite time). Then, again because $\varepsilon > 1$, the dirty input is not essential to final production and therefore, provided that initial environmental quality is sufficiently high, an environmental disaster can be avoided while the economy achieves positive long-run growth at the rate $\gamma\eta_c$. This discussion establishes the following proposition. (Online Appendix II provides a formal proof and also analyzes the case in which $\varepsilon < 1$.)

PROPOSITION 7: *Suppose the two inputs are substitutes* ($\varepsilon > 1$). *Then innovation in the long run will be directed towards the clean sector only, and the economy will grow at rate* $\gamma\eta_c$. *Provided that* $\bar{S}$ *is sufficiently high, an environmental disaster is avoided under laissez-faire.*

The most important result in this proposition is that when the exhaustible resource is necessary for production of the dirty input, the market generates incentives for research to be directed towards the clean sector, and these market-generated incentives may be sufficient for the prevention of an environmental disaster. This contrasts with the result that an environmental disaster is unavoidable under laissez-faire without the exhaustible resource. Therefore, to the extent that in practice the increasing price of oil and the higher costs of oil extraction will create a natural move away from dirty inputs, the implications of growth are not as damaging to the environment as in the baseline case with $\alpha_2 = 0$. Nevertheless, because of the environmental and the knowledge externalities (and also because of the failure to correctly price the resource), the laissez-faire equilibrium is still Pareto suboptimal.

B. *Optimal Environmental Regulation with Exhaustible Resources*

We now briefly discuss the structure of optimal policy in the presence of the exhaustible resource. The socially optimal allocation maximizes (1) now subject to the constraints (4), (5), (6), (7), (8), (9), (11), (12) and the resource constraint $Q_t \geq 0$ for all t.

As in Section III, the socially optimal allocation will correct for the monopoly distortions by subsidizing the use of machines in the two sectors and will again introduce a wedge between the shadow price of the dirty input and its marginal product in the production of the final good, equivalent to a tax on dirty input production. In addition, because the private cost of extraction is $c(Q_t)$ (i.e., does not incorporate the scarcity value of the exhaustible resource), the socially optimal allocation will also use a "resource tax" to create a wedge between the cost of extraction and the social value of the exhaustible resource. The next proposition summarizes the structure of optimal policy in this case.

PROPOSITION 8: *The socially optimal allocation can be implemented using a "carbon" tax* (*i.e., a tax on the use of the dirty input*), *a subsidy to clean research, a subsidy on the use of all machines, and a resource tax* (*all proceeds from taxes/subsidies being redistributed/financed lump sum*). *The resource tax must be maintained forever.*

The proof of this proposition is presented in online Appendix II, which also shows that several quantitative features of the optimal policy in this case are similar to the economy without the exhaustible resource.

C. *Equilibrium and Optimal Policy under the Hotelling Rule*

We next investigate the implications of having well-defined property rights to the exhaustible resource vested in price-taking infinitely lived profit-maximizing firms (see Golosov et al. 2009 for a recent treatment of this case). This implies that the price of the exhaustible resource will be determined by the Hotelling rule.[15] In particular, let us suppose for simplicity that the cost of extraction $c(Q_t)$ is constant and equal to $c > 0$. Then the price of the exhaustible resource, P_t, has to be such that the marginal value of one additional unit of extraction today must be equal to

[15] Yet another alternative would be to have the exhaustible resource owned by a single entity (or consortium), which would not only choose its price according to its scarcity but would also attempt to deviate from the Hotelling rule to internalize the environmental externalities. We find this case empirically less relevant and do not focus on it.

the discounted value of an additional unit extracted tomorrow. More formally, the Hotelling rule in this case takes the form

$$\frac{\partial u(C_t, S_t)}{\partial C}(P_t - c) = \frac{1}{1+\rho}\frac{\partial u(C_{t+1}, S_{t+1})}{\partial C}(P_{t+1} - c). \tag{26}$$

We further simplify the analysis by assuming a constant coefficient of relative risk aversion σ in consumption, and separable preferences between consumption and environmental quality:

$$u(C_t, S_t) = \frac{C_t^{1-\sigma}}{1-\sigma} + \nu(S_t), \tag{27}$$

where $\nu' > 0$ and $\nu'' < 0$. Then the Hotelling rule, (26), implies that the price P_t of the resource must asymptotically grow at the interest rate r, given from the consumption Euler as:

$$r = (1+\rho)(1+g)^{\sigma} - 1, \tag{28}$$

where g is the asymptotic growth rate of consumption.

The next proposition shows that relative to the case analyzed in the previous two subsections, avoiding an environmental disaster becomes more difficult when the price of the exhaustible resource is given by the Hotelling rule.

PROPOSITION 9: *If the discount rate* ρ *and the elasticity of substitution* ε *are both sufficiently high* (*in particular, if* $\ln(1+\rho) > (1-\alpha_1)\ln(1+\gamma\max\{\eta_d, \eta_c\})/\alpha_2$, *and* $\varepsilon > 1/(2-\alpha_1-\alpha)$), *then asymptotically innovation occurs in the clean sector only, and a disaster is avoided under laissez-faire provided that the initial environmental quality,* $\overline{S}$, *is sufficiently high. However, if the discount rate and the elasticity of substitution are sufficiently low* (*in particular, if* $\ln(1+\rho) < (1/\varepsilon - (1-\alpha) - \alpha_2\sigma)\ln(1+\gamma\eta_c)/\alpha_2$ *and* $\ln(1+\rho) \neq (1-\alpha_1)\ln(1+\gamma\eta_d)/\alpha_2$), *then a disaster cannot be avoided under laissez-faire.*

PROOF:

See online Appendix II.

Intuitively, if the price of the resource P_t increases more slowly over time than productivity in the dirty sector, A_{dt}, then under laissez-faire, innovation continues to take place in the dirty sector forever and the growth in the production of the dirty input leads to an environmental disaster. This case arises when the discount rate ρ is sufficiently small. An environmental disaster can be avoided only if the price P_t increases sufficiently fast so that in finite time innovation shifts entirely to the clean sector. This in turn requires the discount rate ρ to be sufficiently high. However, for the same reasons as those highlighted in Section II, such a switch is not sufficient to avoid an environmental disaster unless clean and dirty sectors are "strong substitutes," which now corresponds to the case where $\varepsilon > 1/(2-\alpha_1-\alpha)$.

It can also be shown that a temporary research subsidy is now sufficient to avoid a disaster when $\varepsilon > 1/[1-\alpha+\alpha_2(\ln(1+\rho)/\ln(1+\gamma\eta_c)+\sigma)]$. This threshold is lower than the corresponding threshold $1/(1 - \alpha)$ in the case without the exhaustible resource because dirty inputs are now using the exhaustible resource, which has a price growing at the rate $(1 + \rho)(1 + \gamma\eta_c)^{\sigma} - 1$. This is also the reason why this threshold is decreasing in the share of the exhaustible resource in the production of dirty input. Finally, one can show that the optimal policy is identical to that characterized in Section IVB, except that the resource tax is no longer necessary.

D. *Pollution from Exhaustible Resources*

The introduction of exhaustible resources also enables us to study the case where these are the source of all pollution and environmental degradation. In particular, we could change equation (12) to $S_{t+1} = (1 + \delta)S_t - \xi R_t$. In this case, it can be shown that total environmental damage is bounded above by ξQ_0, which implies that for sufficiently large initial environmental quality S_0 a disaster is always avoided. Nevertheless, the structure of optimal policy is still similar to our baseline model, though it can now be implemented with a subsidy to the use of machines, a subsidy to clean research, and a resource tax, but without a carbon tax, as the resource tax plays the role of a carbon tax in this case. The socially optimal allocation of resources may or may not induce a full switch to clean innovation, but when it does, subsidies to clean research are necessary.[16]

V. An Example

In this section, we report the results of a simple quantitative example. We focus on the economy without exhaustible resources (i.e., $\alpha_2 = 0$).[17] Our objective is not to provide a comprehensive quantitative evaluation but to highlight the effects of different values of the discount rate and the elasticity of substitution on the form of optimal environmental regulation and the resulting timing of a switch (of R&D and production) to clean technology.

A. *Parameter Choices*

We take a period in our model to correspond to five years. We set $\eta_c = \eta_d = 0.02$ (per annum) and $\gamma = 1$ so that the long-run annual growth rate is equal to 2 percent (which matches Nordhaus's assumptions in his 2007 DICE calibration). We take $\alpha = 1/3$ (so that the share of national income spent on machines is approximately equal to the share of capital). We suppose that before the implementation of the optimal policy the carbon tax is 0. To focus on the implications of the environmental externality, we also assume that the subsidy to machines is present throughout. We compute the values of clean and dirty technologies one period before the implementation

[16] In particular, when the utility function is given by (27) and the extraction cost is constant, as in the previous subsection, it can be shown that the optimal policy involves a switch to clean innovation when $\sigma < 1$, ρ is sufficiently small, and $g_d < g_c$, where g_d, defined by $\ln(1 + g_d) = (1 - \alpha_1)\ln(1 + \gamma\eta_d) - \alpha_2 \ln(1 + \rho))/(1 - \alpha + \alpha_2\sigma)$, is the long-run growth rate when innovations take place in the dirty sector only, and $g_c = \gamma\eta_c$ is the long-run growth rate when innovations take place in the clean sector only. The intuition is that when $\sigma < 1$ and ρ is sufficiently small, the social planner prefers the policy alternative that maximizes growth (subject to avoiding environmental disaster), which in this case is the policy inducing a full switch to clean innovation.

[17] The online Appendix II shows that the results are similar in the presence of exhaustible resources.

of the optimal policy, denoted by $A_{c,-1}$ and $A_{d,-1}$, to match the implied values of $Y_{c,-1}$ and $Y_{d,-1}$ to the production of nonfossil and fossil fuel in the world primary energy supply from 2002 to 2006 (according to the Energy Information Administration, 2009 data). Note that in all our exercises, when ε varies, $A_{c,-1}$ and $A_{d,-1}$ also need to be adjusted (in particular, a higher ε leads to a higher ratio of $A_{c,-1}/A_{d,-1}$).

Estimating the economywide elasticity of substitution is beyond the scope of the current article. We simply note that since fossil and nonfossil fuels should be close substitutes (at the very least, once nonfossil fuels can be transported efficiently), reasonable values of ε should be quite high. Here we consider two different values for ε: a low value of $\varepsilon = 3$ and a high value of $\varepsilon = 10$. Contrasting what happens under these two values will allow us to highlight the crucial role of the elasticity of substitution in determining the form of the optimal policy.

To relate the environmental quality variable S to the atmospheric concentration of carbon, we use a common approximation to the relationship between the increase in temperature since preindustrial times (in degrees Celsius), Δ, and the atmospheric concentration of carbon dioxide (CO_2 in ppm):

$$\Delta \simeq 3\log_2(C_{CO2}/280). \tag{29}$$

This equation implies that a doubling of atmospheric concentration in CO_2 leads to a 3°C increase in current temperature (see, e.g., IPCC 2007). We define a disaster as an increase in temperature equal to $\Delta_{disaster} = 6$°C (for example, Stern 2007 reports that increases in temperature of more than 5°C, which among other things will lead to the melting of the Greenland Ice Sheet, significantly raising sea levels, are likely to generate "catastrophic" outcomes including major economic and social disruptions and large-scale population movements). Equation (29) then yields the corresponding disaster level of CO_2 concentration, $C_{CO2,disaster}$, and we set $S = C_{CO2,disaster} - \max\{C_{CO_2}, 280\}$. We also relax the assumption that $S_0 = \overline{S}$ and set the initial environmental quality S_0 to correspond to the current atmospheric concentration of 379 ppm.

We estimate parameter ξ from the observed value of Y_d and the annual emission of CO_2 (ξY_d in our model) between 2002 and 2006, and choose δ such that only half of the amount of emitted carbon contributes to increasing CO_2 concentration in the atmosphere (the rest being offset by "environmental regeneration," see IPCC 2007, 2008).

Nordhaus—and much of the literature following his work—assumes that environmental quality affects aggregate productivity. Instead, we formulated our model under the assumption that environmental quality directly affects utility. To highlight the similarities and the differences between our model and existing quantitative models with exogenous technology, we choose the parameters such that the welfare consequences of changes in temperature (for the range of changes observed so far) are the same in our model as in previous work. We parameterize the utility function as

$$u(C_t, S_t) = \frac{(\phi(S_t)\,C_t)^{1-\sigma}}{1-\sigma}, \tag{30}$$

with $\sigma = 2$, which matches Nordhaus's choice of intertemporal elasticity of substitution. In addition, this utility function contains the term $\phi(S)$ for the costs from the degradation of environmental quality. We choose this function as

$$(31)\,\phi(S) = \varphi(\Delta(S)) \equiv \frac{(\Delta_{disaster} - \Delta(S))^{\lambda} - \lambda \Delta_{disaster}^{\lambda-1}(\Delta_{disaster} - \Delta(S))}{(1 - \lambda)\Delta_{disaster}^{\lambda}},$$

which satisfies our assumptions (2) and (3) above. Matching this function with Nordhaus's damage function over the range of temperature increases up to 3°C leads to a value of $\lambda = 0.1443$.

The debate between Stern and Nordhaus highlighted the importance of the discount rate when determining the optimal environmental policy. In the following simulations we consider two different values for the discount rate: the Stern discount rate of 0.001 per annum (which we write as $\rho = 0.001$), and the Nordhaus discount rate of 0.015 per annum ($\rho = 0.015$, which, as in Nordhaus, corresponds to an annual long-run interest rate of about $r = \rho + \sigma g = 5.5$ percent).

B. *Results*

Figure 1 shows the subsidy to the clean sector, the allocation of scientists to clean technologies, the "carbon" tax, the share of clean inputs in total production, and the increase in temperature in the optimal allocation for the following configurations: $[\varepsilon = 10,\ \rho = 0.015]$, $[\varepsilon = 3,\ \rho = 0.001]$, and $[\varepsilon = 3,\ \rho = 0.015]$. The choice of $[\varepsilon = 10,\ \rho = 0.001]$ leads to identical results to those obtained from $[\varepsilon = 10,\ \rho = 0.015]$ and is not shown to make the figure easier to read.

Figure 1, panel B shows that when $\varepsilon = 10$ or when $\varepsilon = 3$ and $\rho = 0.001$, the optimal policy involves an immediate switch of all research activities towards clean technologies. When $\varepsilon = 3$ and $\rho = 0.015$, the switch toward clean research occurs around year 50. As shown in Figure 1, panel A, the optimal subsidy to clean research is temporary, and it is lower and of shorter duration when $\varepsilon = 10$, because in this case the initial gap between clean and dirty technologies consistent with the observed share of dirty inputs is smaller. When $\varepsilon = 3$, the optimal subsidy is larger and lasts longer, particularly when $\rho = 0.015$, because in this case the switch to clean research occurs later.

Figure 1, panel C shows that when $\varepsilon = 10$, the carbon tax is very low and applies only for a limited period because the rapid switch to clean inputs makes this tax unnecessary. In contrast, when $\varepsilon = 3$ and $\rho = 0.015$, because the switch of both innovation and production to the clean sector is delayed, there is a much higher and initially (for over 185 years) increasing carbon tax. Figure 1, panel D shows that when $\varepsilon = 10$, the clean sector takes over most of input production quite rapidly (it takes only 30 years for 90 percent of input production to switch to the clean sector). In contrast, when $\varepsilon = 3$ and $\rho = 0.001$, even though the switch to clean research is immediate, it takes much longer (over 100 years) for 90 percent of inputs to be supplied by the clean sector. Figure 1, panel E shows that when $\varepsilon = 10$, there is a small increase, followed by a decrease, in temperature (going back to its preindustrial level after about 90 years). The pattern is similar, though the increase and the subsequent decline are more protracted when $\varepsilon = 3$ and $\rho = 0.001$. Finally, when $\varepsilon = 3$ and $\rho = 0.015$, temperature

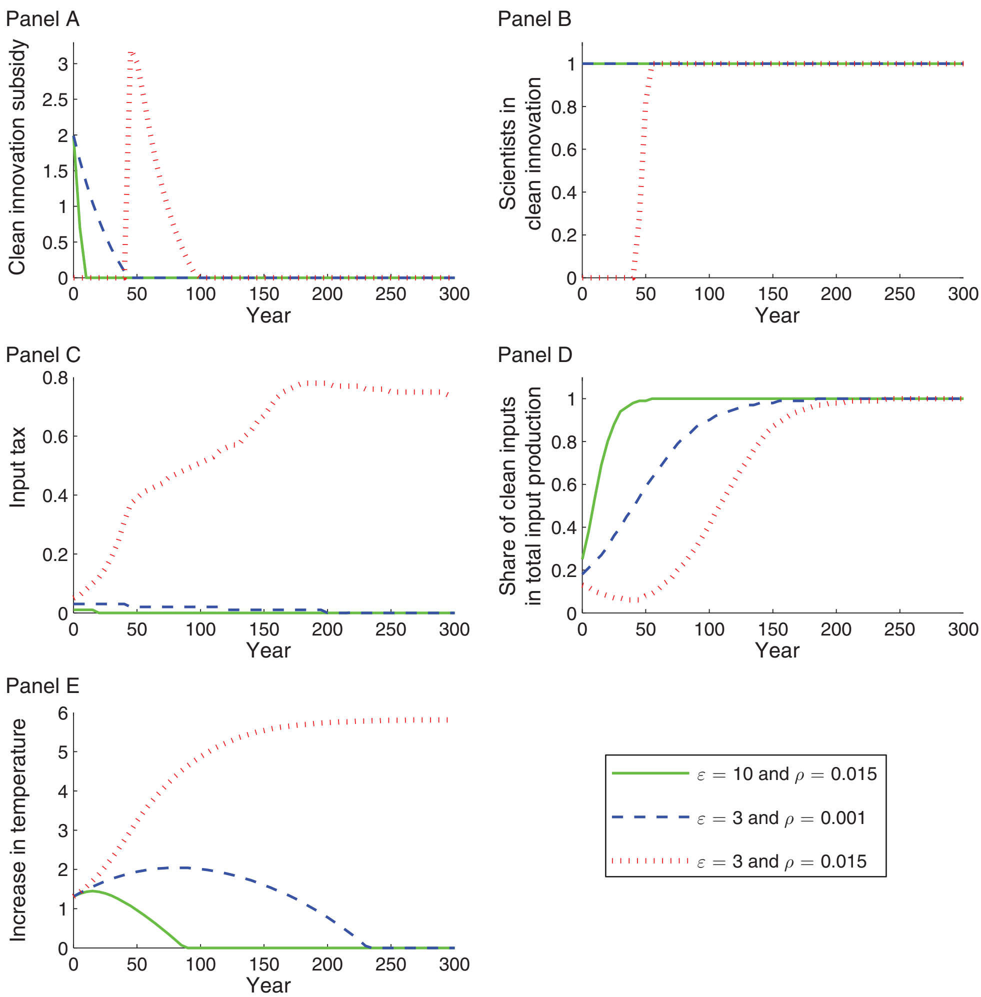

Figure 1. Optimal Environmental Policy for Different Values of ε and ρ

keeps increasing for about 300 years before reaching a maximum fairly close to the disaster level. Overall, these results suggest that if the elasticity of substitution between clean and dirty inputs is sufficiently high, then whether one uses the Nordhaus or the Stern discount rate has little bearing on the nature of the optimal environmental policy.

Corollary 1 in Section IIC related the costs of delayed intervention to the number of additional periods of slow growth that such a delay would induce. Table 1 here shows the welfare costs of delaying the implementation of the optimal policy (i.e., of maintaining the clean innovation subsidy and the carbon tax at zero for a while before implementing the optimal policy) for different values of ε and ρ.[18] Welfare costs are measured as the equivalent percentage reduction in per period consumption relative to the allocation with immediate intervention (we assume that when intervention starts, it takes the optimal form). The table shows that delay costs can be substantial. For example, with $\varepsilon = 10$ and $\rho = 0.001$, a ten-year delay is

[18] The optimal subsidy on machines is maintained during the period of delay.

Table 1—Welfare Costs of Delayed Intervention as a Function of the Elasticity of Substitution and the Discount Rate

Elasticity of substitution ε	10		3	
Discount rate ρ	0.001	0.015	0.001	0.015
delay = 10 years	8.50	0.69	2.30	0.04
delay = 20 years	13.37	0.73	3.88	0.10
delay = 30 years	16.49	0.79	5.52	0.20

Note: Percentage reductions in consumption relative to immediate intervention.

equivalent to an 8.50 percent decline in consumption. Moreover, the cost of delay increases with the duration of the delay and the elasticity of substitution between the two inputs. Intuitively, the latter result arises because when the two inputs are close substitutes, further advances in the dirty technology that occur before the optimal policy is implemented do not contribute much to aggregate output once the switch to clean research and production takes place. The cost of delay also decreases with the discount rate because the benefit from delaying intervention, due to higher consumption early on, increases with the discount rate.

Finally, we briefly discuss the welfare costs of relying solely on a carbon (input) tax instead of combining it with the subsidy to clean research (i.e., the "second-best" instead of "first-best" derived in Proposition 6). Without the subsidy to clean research, the carbon tax needs to be significantly higher. For example, when $\varepsilon = 10$ and $\rho = 0.015$, the initial value of the carbon tax in the second-best needs to be 40 times higher than in the first-best. The higher tax level creates a greater reduction in production and consumption in the short run. Table 2 shows that the welfare loss in the second-best relative to the first-best can be significant (though it is typically smaller than the costs of delay shown in Table 1). It is smaller when the elasticity of substitution is high, since in this case a relatively small carbon tax is sufficient to redirect R&D towards clean technologies; and it is greater when the discount rate is high, because a higher discount rate puts greater weight on earlier periods where a significantly higher carbon tax needs to be imposed in the second-best.

VI. Conclusion

In this article we introduced endogenous and directed technical change in a growth model with environmental constraints and limited resources. We characterized the structure of equilibria and the dynamic tax/subsidy policies that achieve sustainable growth or maximize intertemporal welfare. The long-run properties of both the laissez-faire equilibrium and the social optimum (or the necessary policies to avoid environmental disaster) are related to the degree of substitutability between clean and dirty inputs, to whether dirty input production uses exhaustible resources, and to initial environmental and resource stocks.

The main implications of factoring in the importance of directed technical change are as follows: (i) when the inputs are sufficiently substitutable, sustainable long-run growth can be achieved using temporary policy intervention (e.g., a temporary research subsidy to the clean sector), and need not involve long-run distortions; (ii) optimal

Table 2—Welfare Costs of Relying Solely on Carbon Tax as a Function of the Elasticity of Substitution and the Discount Rate

Elasticity of substitution ε	10		3	
Discount rate ρ	0.001	0.015	0.001	0.015
Welfare cost	1.02	1.66	1.92	3.15

Note: Percentage reductions in consumption relative to immediate intervention.

policy involves both carbon taxes and research subsidies, so that excessive use of carbon taxes can be avoided; (iii) delay in intervention is costly: the sooner and the stronger the policy response, the shorter will the slow growth transition phase be; (iv) the use of an exhaustible resource in dirty input production helps the switch to clean innovation under laissez-faire. Thus the response of technology to policy leads to a more optimistic scenario than what emerges from models with exogenous technology. However, directed technical change also calls for immediate and decisive action in contrast to the implications of several exogenous technology models used in previous economic analyses.

A simple quantitative evaluation suggests that, provided that the elasticity of substitution between clean and dirty inputs is sufficiently high, optimal environmental regulation should involve an immediate switch of R&D resources to clean technology, followed by a gradual switch of all production to clean inputs. This conclusion appears robust to the range of discount rates used in the Stern report and in Nordhaus's work (which lead to very different policy conclusions in models with exogenous technology). Interestingly, in most cases, optimal environmental regulation involves small carbon taxes because research subsidies are able to redirect innovation to clean technologies before there is more extensive environmental damage.

Our paper is a first step toward a comprehensive framework that can be used for theoretical and quantitative analysis of environmental regulation with endogenous technology. Several directions of future research appear fruitful. First, it would be useful to develop a multicountry model with endogenous technology and environmental constraints, which can be used to discuss issues of global policy coordination and the degree to which international trade should be linked to environmental policies. Second, an interesting direction is to incorporate "environmental risk" into this framework, for example, because of the ex ante uncertainty on the regeneration rate, δ, or on future costs of environmental damage. Another line of important future research would be to exploit macroeconomic and microeconomic (firm- and industry-level) data to estimate the relevant elasticity of substitution between clean and dirty inputs.

Appendix I

A. *Solving for the Laissez-Faire Equilibrium*

In this Appendix we solve for the profit-maximization of machine producers and express the price and labor allocation ratios as functions of the relative aggregate productivities of clean and dirty technologies in the laissez-faire equilibrium.

The profit-maximization problem of the producer of machine i at time t in sector $j \in \{c, d\}$ can be written as

$$\max_{x_{jit}, L_{jt}} \left\{ p_{jt} L_{jt}^{1-\alpha} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} \; di \; - \; w_t L_{jt} \; - \; \int_0^1 p_{jit} x_{jit} \; di \right\},$$

and leads to the following iso-elastic inverse demand curve:

$$x_{jit} \; = \; \left(\frac{\alpha p_{jt}}{p_{jit}} \right)^{\frac{1}{1-\alpha}} A_{jit} L_{jt}. \tag{A.1}$$

The monopolist producer of machine i in sector j chooses p_{jit} and x_{jit} to maximize profits $\pi_{jit} = (p_{jit} - \psi) x_{jit}$, subject to the inverse demand curve (A.1). Given this isoelastic demand, the profit-maximizing price is a constant markup over marginal cost, thus $p_{jit} = \psi / \alpha$. Recalling the normalization $\psi \equiv \alpha^2$, this implies that $p_{jit} = \alpha$ and thus the equilibrium demand for machines i in sector j is obtained as

$$x_{jit} \; = \; p_{jt}^{\frac{1}{1-\alpha}} L_{jt} A_{jit}. \tag{A.2}$$

Equilibrium profits for the monopolist are then given by (15) in the text.

Next, combining equation (A.2) with the first-order condition with respect to labor, $(1-\alpha) p_{jt} L_{jt}^{-\alpha} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di = w_t$, and using (10) gives the relative prices of clean and dirty inputs as

$$\frac{p_{ct}}{p_{dt}} \; = \; \left(\frac{A_{ct}}{A_{dt}} \right)^{-(1-\alpha)}. \tag{A.3}$$

This equation formalizes the natural idea that the input produced with more productive machines will be relatively cheaper.

Equation (A.2) together with (5) gives the equilibrium production level of input j as

$$Y_{jt} \; = \; \left(p_{jt} \right)^{\frac{\alpha}{1-\alpha}} A_{jt} L_{jt}. \tag{A.4}$$

Combining (A.4) with (13), then using (A.3) and the definition of $\varphi \equiv (1-\alpha) \times (1-\varepsilon)$, we obtain the relationship between relative productivities and relative employment as

$$\frac{L_{ct}}{L_{dt}} \; = \; \left(\frac{p_{ct}}{p_{dt}} \right)^{-\frac{\varphi - 1}{1-\alpha}} \frac{A_{dt}}{A_{ct}} \; = \; \left(\frac{A_{ct}}{A_{dt}} \right)^{-\varphi}. \tag{A.5}$$

Finally, combining (A.3) and (A.5) with (17) gives (18) in the text.

B. *Equilibrium Allocations of Scientists*

We now characterize the equilibrium allocation of innovation effort across the two sectors for any value of the elasticity parameter ε and provide a proof of Lemma 1. Defining

$$f(s) \equiv \frac{\eta_c}{\eta_d}\left(\frac{1 + \gamma\eta_c s}{1 + \gamma\eta_d(1 - s)}\right)^{-\varphi-1}\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{-\varphi},$$

for $s \in [0,1]$, we can rewrite (18) as $\Pi_{ct}/\Pi_{dt} = f(s_{ct})$. Clearly, if $f(1) > 1$, then $s = 1$ is an equilibrium; if $f(0) < 1$, then $s = 0$ is an equilibrium; and finally if $f(s^*) = 1$ for some $s^* \in (0,1)$, then s^* is an equilibrium. Given these observations, we have:

1. If $1 + \varphi > 0$ (or equivalently $\varepsilon < (2 - \alpha)/(1 - \alpha)$), then $f(s)$ is strictly decreasing in s. Then it immediately follows that: (i) if $f(1) > 1$, then $s = 1$ is the unique equilibrium (we only have a corner solution in that case); (ii) if $f(0) < 1$, then $s = 0$ is the unique equilibrium (again a corner solution); (iii) if $f(0) > 1 > f(1)$, then by continuity there exists a unique $s^* \in (0,1)$ such that $f(s^*) = 1$, which is the unique (interior) equilibrium.

2. If $1 + \varphi < 0$ (or equivalently $\varepsilon > (2 - \alpha)/(1 - \alpha)$), then $f(s)$ is strictly increasing in s. In that case: (i) if $1 < f(0) < f(1)$, then $s = 1$ is the unique equilibrium; (ii) if $f(0) < f(1) < 1$, then $s = 0$ is the unique equilibrium; (iii) if $f(0) < 1 < f(1)$, then there are three equilibria, an interior one $s = s^* \in (0,1)$ where s^* is such that $f(s^*) = 1$, $s = 0$ and $s = 1$.

3. If $1 + \varphi = 0$, then $f(s) \equiv f$ is a constant. If f is greater than 1, then $s = 1$ is the unique equilibrium; if it is less than one, then $s = 0$ is the unique equilibrium.

This characterizes the allocation of scientists and implies the results in Lemma 1.

C. *Proof of Proposition 1*

Assumption 1 together with the characterization of equilibrium allocation of scientists above implies that, initially, innovation takes place in the dirty sector only ($s_{dt} = 1$ and $s_{ct} = 0$). From (11), this widens the gap between clean and dirty technologies and ensures that $s_{dt+1} = 1$ and $s_{ct+1} = 0$, and so on in subsequent periods. This shows that under Assumption 1, the equilibrium is uniquely defined under laissez-faire and involves $s_{dt} = 1$ and $s_{ct} = 0$ for all t.

D. *Proof of Proposition 5*

Let λ_t denote the Lagrange multiplier for (4), which is naturally also the shadow value of one unit of final good production. The first-order condition with respect to Y_t implies that this shadow value is equal to the Lagrange multiplier for (8), so that it is also equal to the shadow value of one unit of consumption. Then the first-order condition with respect to C_t yields

$$\lambda_t = \frac{1}{(1+\rho)^t}\frac{\partial u(C_t, S_t)}{\partial C}, \tag{A.6}$$

so that the shadow value of the final good is equal to the marginal utility of consumption.

Next, letting ω_t denote the Lagrange multiplier for the environmental equation (12), the first-order condition with respect to S_t gives

$$\omega_t = \frac{1}{(1+\rho)^t}\frac{\partial u(C_t, S_t)}{\partial S} + (1+\delta) I_{S_t<\bar{S}}\,\omega_{t+1}, \tag{A.7}$$

where $I_{S_t<\bar{S}}$ is equal to 1 if $S_t < \bar{S}$ and to 0 otherwise. This implies that the shadow value of environmental quality at time t is equal to the marginal utility that it generates in this period plus the shadow value of $(1+\delta)$ units of environmental quality at time $t+1$ (as one unit of environmental quality at time t generates $1+\delta$ units at time $t+1$). Solving (A.7) recursively, we obtain that the shadow value of environmental quality at time t is

$$\omega_t = \sum_{v=t}^{\infty}(1+\delta)^{v-t}\frac{1}{(1+\rho)^v}\frac{\partial u(C_v, S_v)}{\partial S} I_{S_t,\ldots,S_v<\bar{S}}, \tag{A.8}$$

where $I_{S_t,\ldots,S_v<\bar{S}}$ takes value 1 if $S_t,\ldots,S_\nu < \bar{S}$ and 0 otherwise. Given the assumption that $\partial u(C,\bar{S})/\partial S = 0$, this equation also implies that if for all $v > T$, $S_v = \bar{S}$, then $\omega_t = 0$ for all $t > T$.

Defining λ_{jt} as the Lagrange multiplier for (5), the ratio λ_{jt}/λ_t can be interpreted as the shadow price of input j at time t (relative to the price of the final good). To emphasize this interpretation, we will denote this ratio by $\hat{p}_{jt}$. The first-order conditions with respect to Y_{ct} and Y_{dt} then give

$$Y_{ct}^{\frac{-1}{\varepsilon}}\left(Y_{ct}^{\frac{\varepsilon-1}{\varepsilon}} + Y_{dt}^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{1}{\varepsilon-1}} = \hat{p}_{ct} \tag{A.9}$$

$$Y_{dt}^{\frac{-1}{\varepsilon}}\left(Y_{ct}^{\frac{\varepsilon-1}{\varepsilon}} + Y_{dt}^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{1}{\varepsilon-1}} - \frac{\omega_{t+1}\xi}{\lambda_t} = \hat{p}_{dt}.$$

These equations imply that compared to the laissez-faire equilibrium, the social planner introduces a wedge of $\omega_{t+1}\xi/\lambda_t$ between the marginal product of the dirty input and its price. This wedge $\omega_{t+1}\xi/\lambda_t$ is equal to the environmental cost of an additional unit of the dirty input (evaluated in terms of units of the final good at time t; recall that one unit of dirty production at time t destroys ξ units of environmental quality at time $t+1$). Naturally, this wedge is also equivalent to a tax of

$$\tau_t = \frac{\omega_{t+1}\xi}{\lambda_t\hat{p}_{dt}} \tag{A.10}$$

on the use of dirty input by the final good producer. This tax rate will be higher when the shadow value of environmental quality is greater; when the marginal utility of consumption today is lower; and when the price of dirty input is lower. Plugging (A.8) and (A.6) in (A.10) we get (23).

Next, the subsidy to the use of all machines can be derived from the first-order condition with respect to η_{ji}:

$$x_{jit} = \left(\frac{\alpha}{\psi}\hat{p}_{jt}\right)^{1/(1-\alpha)} A_{jit} L_{jt}. \tag{A.11}$$

Comparing this expression to the equilibrium inverse demand, (A.1) highlights that existing machines will be used more intensively in the socially planned allocation. This is a natural consequence of the monopoly distortions and can also be interpreted as the socially planned allocation involving a subsidy of $1-\alpha$ in the use of machines, so that their price should be identical to the marginal cost, i.e., $(1-(1-\alpha))\psi/\alpha = \psi \equiv \alpha^2$.

We can combine (A.11) with (5) to obtain

$$Y_{jt} = \left(\frac{\alpha}{\psi}\hat{p}_{jt}\right)^{\alpha/(1-\alpha)} A_{jt} L_{jt}, \tag{A.12}$$

so that for given price, average technology and labor allocation, the production of each input is scaled up by a factor $\alpha^{-\alpha/(1-\alpha)}$ compared to the laissez-faire equilibrium (this results from the more intensive use of machines in the socially planned allocation).

Finally, the socially optimal allocation must correct for the knowledge externality. Let μ_{jt} denote the Lagrange multiplier for equation (11) for $j = c, d$ (corresponding to the shadow value of average productivity in sector j at time t). The relevant first-order condition gives

$$\mu_{jt} = \lambda_t\left(\frac{\alpha}{\psi}\right)^{\alpha/(1-\alpha)} (1-\alpha)\hat{p}_{jt}^{1/(1-\alpha)} L_{jt} + (1+\gamma\eta_j s_{jt+1})\mu_{jt+1}. \tag{A.13}$$

Intuitively, the shadow value of a unit increase in average productivity in sector $j \in \{c, d\}$ is equal to its marginal contribution to time-t utility plus its shadow value at time $t+1$ times $(1+\gamma\eta_j s_{jt+1})$ (the further productivity increase it enables at time $t+1$). This last term captures the intertemporal knowledge externality.

In the optimal allocation of resources, scientists will be allocated towards the sector with the higher social gain from innovation, as measured by $\gamma\eta_j\mu_{jt}A_{jt-1}$. Using (A.13), we then have that the social planner will allocate scientists to the clean sector whenever the ratio

$$\frac{\eta_c(1+\gamma\eta_c s_{ct})^{-1}\sum_{v\geq t}\lambda_v\hat{p}_{cv}^{1/(1-\alpha)}L_{cv}A_{cv}}{\eta_d(1+\gamma\eta_d s_{dt})^{-1}\sum_{v\geq t}\lambda_v\hat{p}_{dv}^{1/(1-\alpha)}L_{dv}A_{dv}} \tag{A.14}$$

is greater than 1 (combining (A.6) and (A.14) we obtain (24)). The social planner can implement this optimal allocation through a subsidy q_t to clean research. To determine this subsidy, first note that in the optimal allocation the shadow values of the clean and dirty inputs satisfy

$$\hat{p}_{ct}^{1/(1-\alpha)} A_{ct} = \hat{p}_{dt}^{1/(1-\alpha)} A_{dt}. \tag{A.15}$$

Then, using (A.9), (A.12), and (A.15), we obtain

$$\frac{L_{ct}}{L_{dt}} = (1 + \tau_t)^{\varepsilon} \left(\frac{A_{ct}}{A_{dt}}\right)^{-\varphi}. \tag{A.16}$$

Next, using (A.11), pretax profits are $\pi_{jit} = (1-\alpha)(\alpha/\psi)^{\alpha/(1-\alpha)} \hat{p}_{jt}^{1/(1-\alpha)} A_{jit} L_{jt}$. Therefore, for given subsidy q_t, the ratio of expected profits from innovation in sectors c and d, the equivalent of (18) in the text, can be written as

$$\frac{\Pi_{ct}}{\Pi_{dt}} = (1 + q_t)\frac{\eta_c}{\eta_d}\left(\frac{1 + \gamma\eta_c s_{ct}}{1 + \gamma\eta_d s_{dt}}\right)^{-\varphi-1} (1 + \tau_t)^{\varepsilon}\left(\frac{A_{ct-1}}{A_{dt-1}}\right)^{-\varphi}. \tag{A.17}$$

Clearly, when the optimal allocation involves $s_{ct} = 1$, we can choose q_t to make this expression greater than one. Or, more explicitly, we can set

$$q_t \geq \hat{q}_t \equiv \frac{\eta_d}{\eta_c}(1 + \gamma\eta_d)^{-\varphi-1}(1 + \tau_t)^{-\varepsilon}\left(\frac{A_{dt-1}}{A_{ct-1}}\right)^{-\varphi} - 1.$$

When the optimal allocation involves $s_{ct} \in (0,1)$, then setting q_t to ensure that $\Pi_{ct}/\Pi_{dt} = 1$ achieves the desired objective.

REFERENCES

Acemoglu, Daron. 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics* 113(4): 1055–89.

Acemoglu, Daron. 2002. "Directed Technical Change." *Review of Economic Studies* 69(4): 781–809.

Acemoglu, Daron. 2007. "Equilibrium Bias of Technology." *Econometrica* 75(5): 1371–1409.

Acemoglu, Daron. 2009. *Introduction to Modern Economic Growth*. New York: Princeton University Press.

Acemoglu, Daron. 2010. "When Does Labor Scarcity Encourage Innovation?" *The Journal of Political Economy* 118(6): 1037–78 .

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.1.131.

Aghion, Philippe, and Peter Howitt. 1998. *Endogenous Growth Theory*. "Problems and Solutions." Cambridge, MA: MIT Press.

Aghion, Philippe, and Peter Howitt. 2009. *The Economics of Growth*. With the collaboration of Leonardo Bursztyn. Cambridge, MA: MIT Press.

Bosetti, Valentina, Carlo Carraro, Emanuele Massetti, and Massimo Tavoni. 2008. "International Energy R&D Spillovers and the Economics of Greenhouse Gas Atmospheric Stabilization." *Energy Economics* 30(6): 2912–29.

Bovenberg, A. Lans, and Sjak A. Smulders. 1995. "Environmental Quality and Pollution-Augmenting Technological Change in a Two-Sector Endogenous Growth Model." *Journal of Public Economics* 57(3): 369–91.

Bovenberg, A. Lans, and Sjak A. Smulders. 1996. "Transitional Impacts of Environmental Policy in an Endogenous Growth Model." *International Economic Review* 37(4): 861–93.

Buonanno, Paolo, Carlo Carraro, and Marzio Galeotti. 2003. "Endogenous Induced Technical Change and the Costs of Kyoto." *Resource and Energy Economics* 25(1): 11–34.

Dasgupta, Partha. 2007. "Commentary: The Stern Review's Economics of Climate Change." *National Institute Economic Review* (199): 4–7.

Dasgupta, Partha. 2008. "Discounting Climate Change." *Journal of Risk and Uncertainty* 37(2–3): 141–69.

Di Maria, Corrado, and Simone Valente. 2006. "The Direction of Technical Change in Capital-Resource Economies." Swiss Federal Institute of Technology-Zurich Working Paper 06/50.

Emanuel, Kerry. 2005. "Increasing Destructiveness of Tropical Cyclones Over the Past 30 Years." *Nature* 436: 686–88.

Energy Information Administration. 2009. "Annual Energy Review, 2008." www.eia.gov/aer.

Fischer, Carolyn, and Richard G. Newell. 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management* 55(2): 142–62.

Gans, Joshua S. 2009. "Innovation and Climate Change Policy." Unpublished.

Gerlagh, Reyer, Snorre Kverndokk, and Knut Einar Rosendahl. 2009. "Optimal Timing of Climate Change Policy: Interaction between Carbon Taxes and Innovation Externalities." *Environmental and Resource Economics* 43(3): 369–90.

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski. 2009. "Optimal Taxes on Fossil Fuel in General Equilibrium." Unpublished.

Goulder, Lawrence H., and Koshy Mathai. 2000. "Optimal CO2 Abatement in the Presence of Induced Technological Change." *Journal of Environmental Economics and Management* 39(1): 1–38.

Goulder, Lawrence H., and Stephen H. Schneider. 1999. "Induced Technological Change and the Attractiveness of CO2 Abatement Policies." *Resource and Energy Economics* 21(3–4): 211–53.

Grimaud, André, and Luc Rouge. 2008. "Environment, Directed Technical Change and Economic Policy." *Environmental and Resource Economics* 41(4): 439–63.

Grübler, Arnulf, and Sabine Messner. 1998. "Technological Change and the Timing of Mitigation Measures." *Energy Economics* 20(5–6): 495–512.

Hepburn, Cameron. 2006. "Regulation by Prices, Quantities, or Both: A Review of Instrument Choice." *Oxford Review of Economic Policy* 22(2): 226–47.

Intergovernmental Panel on Climate Change. 2007. "Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change." In *Climate Change 2007: Synthesis Report*, edited by Rajendra K. Pachauri and Andy Reisinger. Geneva, Switzerland:Intergovernmental Panel on Climate Change.

Jones, Charles I. 2009. "The Costs of Economic Growth." Unpublished.

Landsea, Christopher W. 2005. "Meteorology: Hurricanes and Global Warming." *Nature* 438: E11–E12.

MacCracken, Christopher N., James A. Edmonds, Son H. Kim, and Ronald D. Sands. 1999. "The Economics of the Kyoto Protocol." *The Energy Journal* 20 (Special Issue): 25–72.

Manne, Alan, and Richard Richels. 2004. "The Impact of Learning-by-Doing on the Timing and Costs of CO2 Abatement." *Energy Economics* 26(4): 603–19.

Massetti, Emanuele, Carlo Carraro, and Lea Nicita. 2009. "How Does Climate Policy Affect Technical Change? An Analysis of the Direction and Pace of Technical Progress in a Climate-Economy Model." *The Energy Journal* 30(2): 7–38.

Mendelsohn, Robert O. 2006. "A Critique of the Stern Report." *Regulation* 29(4): 42–46.

Messner, Sabine. 1997. "Endogenized Technological Learning in an Energy Systems Model." *Journal of Evolutionary Economics* 7(3): 291–313.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins. 1999. "The Induced Innovation Hypothesis and Energy-Saving Technological Change." *Quarterly Journal of Economics* 114(3): 941–75.

Nicholls, Robert J., and Jason A. Lowe. 2006. "Climate Stabilisation and Impacts of Sea-level Rise." In *Avoiding Dangerous Climate Change*, edited by Hans J. Schellnhuber, Wolfgang Cramer, Nebojsa Nakicenovic, Tom Wigley and Gary Yohe, 195–201. Cambridge, UK: Cambridge University Press.

Nordhaus, William D. 1994. *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA: MIT Press.

Nordhaus, William D. 2002. "Modeling Induced Innovation in Climate-Change Policy." In *Technological Change and the Environment*, edited by A. Grübler, N. Nakicenovic and W. D. Nordhaus, 182–209. Washington, D.C.: Resources for the Future Press.

Nordhaus, William D. 2007. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature* 45(3): 686–702.

Nordhaus, William D. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies.* New Haven, CT: Yale University Press.

Nordhaus, William D., and Joseph Boyer. 2000. *Warming the World: Economic Models of Global Warming*. Cambridge, MA: MIT Press.

Pizer, William A. 2002. "Combining Price and Quantity Controls to Mitigate Global Climate Change." *Journal of Public Economics* 85(3): 409–34.

Popp, David. 2002. "Induced Innovation and Energy Prices." *American Economic Review* 92(1): 160–80.

Popp, David. 2004. "ENTICE: Endogenous Technological Change in the DICE Model of Global Warming." *Journal of Environmental Economics and Management* 48(1): 742–68.

Stern, Nicholas. 2009. *A Blueprint for a Safer Planet: How to Manage Climate Change and Create a New Era of Progress and Prosperity.* London: Bodley Head.

Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review.* New York: Cambridge University Press.

Stokey, Nancy L. 1998. "Are There Limits to Growth?" *International Economic Review* 39(1): 1–31.

Stott, Peter A., Dáithí A. Stone, and Myles R. Allen. 2004. "Human Contribution to the European Heatwave of 2003." *Nature* 432: 610–614.

Sue Wing, Ian. 2003. "Induced Technical Change and the Cost of Climate Policy." MIT Joint Program on the Science and Policy of Global Change.

Tol, Richard S. J., and Gary W. Yohe. 2006. "A Review of the Stern Review." *World Economics* 7(4): 233–50.

van der Zwaan, B. C. C., R. Gerlagh, G. Klaassen, and L. Schrattenholzer. 2002. "Endogenous Technological Change in Climate Change Modelling." *Energy Economics* 24(1): 1–19.

von Below, David, and Torsten Persson. 2008. "Uncertainty, Climate Change and the Global Economy." Unpublished.

Weitzman, Martin L. 1974. "Prices vs. Quantities." *Review of Economic Studies* 41(4): 477–91.

Weitzman, Martin L. 2007. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature* 45(3): 703–24.

Weitzman, Martin L. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91(1): 1–19.

This article has been cited by:

1. David Autor. 2012. The Journal of Economic Perspectives at 100 (Issues). *Journal of Economic Perspectives* **26**:2, 3-18. [Abstract] [View PDF article] [PDF with links]