Induced Technological Change: Firm Innovatory Responses to Environmental Regulation

This draft: 30 June 2006

Ian Sue Wing [1]

Kennedy School of Government, Harvard University

Joint Program on the Science & Policy of Global Change, MIT

Center for Energy & Environmental Studies and Dept. of Geography & Environment, Boston University

Abstract

In Hicks's (1932) articulation of the induced technical change hypothesis, a change in relative prices stimulates innovation to conserve on relatively expensive inputs. We investigate the workings of this process when price changes result from environmental tax and the regulated firms perform the innovation themselves. We develop a simple dynamic model of a firm which faces a downward-sloping demand curve, produces output using clean and dirty inputs, and invests in clean and dirty research. The tax raises the relative price of the dirty input, increasing the relative attractiveness of pollution-augmenting R&D while crowding out clean R&D. We demonstrate how this effect arises out of the tension between the incentive to innovate to increase revenue and the cost of the research necessary to generate inventions, elucidate its sensitivity to the characteristics of the firm, its market environment and the stringency of the tax, and elaborate its consequences for the firm's profit and pollutant emissions.

Key words: innovation, pollution, dirty inputs, crowding out, Porter Hypothesis
JEL Codes: Q55, O30, D21

Email address: `isw@bu.edu` (Ian Sue Wing).
[1] Address: Rm. 141, 675 Commonwealth Ave., Boston, MA 02215. Phone: (617) 353-4751. Fax: (617) 353-5986. This research was supported by U.S. Department of Energy Office of Science (BER) Grant No. DE-FG02-02ER63484, and by a Harvard Kennedy School REPSOL-YPF Energy Fellowship. This paper has benefited from helpful comments and suggestions by Gib Metcalf, John Reilly, Bill Hogan, David Popp, Larry Goulder, Rob Williams, and seminar participants at the Ohio State University, Harvard University and the Technological Change and the Environment Workshop at Dartmouth College. Errors and omissions remain my own.

1 Introduction

Technological change is one of the most important and least well understood influences on the cost of environmental regulation. There has been intense interest in induced technological change (ITC), the process by which regulatory constraints alter the rate and direction of innovation. [2] A particularly controversial aspect of ITC is the Porter Hypothesis, which posits that improvements in technology induced by mandates to reduce pollution not only mitigate firms' costs of abatement but actually cause their profits to increase (e.g., Ashford et al. 1985; Ashford 1994; Porter and van der Linde 1995). Although this fortuitous outcome has been dismissed as an implausible free lunch (Palmer et al., 1995), the Porterian conjecture nonetheless raises the key questions of what are the precise mechanisms by which environmental regulations influence technological change, and how these influence firms' pollution and profits. The present study provides answers by elucidating how the input price change that results from regulation alters a firm's propensity to invest in different lines of research.

In neoclassical models with perfect markets and no uncertainty, environmental policy constraints have been shown to increase firms' profits only if two conditions are met: (a) innovation improves productivity while simultaneously reducing pollution,and (b) there is some additional market failure or source of increasing returns which prevented the firm from making investments to reap productivity gains in the absence of regulation. These kinds of assumptions are common in papers which investigate firms' decisions to adopt new technology, [3] or their propensity to innovate in response to regulatory constraints (Goulder and Matthai, 2000; Parry et al., 2003). [4] However, the optimistic assumption of com-

[2] The initial articulation of the induced innovation hypothesis is customarily attributed to Hicks (1932, p. 124): "a change in the relative prices of factors of production is itself a spur to invention, and to invention of a particular kind-directed to economizing the use of a factor which has become relatively expensive". For surveys of the early literature on ITC see Binswanger and Ruttan (1978) and Thirtle and Ruttan (1987). Kamien and Schwartz (1968) made the most progress toward a fully-articulated theory of induced innovation, while Magat (1976) provided the first application of ITC to pollution control.
[3] The assumption is that new technologies are more costly than existing capital, but are both cleaner and more efficient. Xepapadeas and DeZeeuw (1999) and Feichtinger et al. (2005) show that in the absence of increasing returns environmental regulation accelerates the obsolescence of old capital, which increases the efficiency of production and mitigates—but does not completely offset—the loss suffered by the regulated firm. Mohr (2002) shows that if there are external economies of scale in adoption, environmental regulation acts to coordinate adoption among firms, thereby increasing profits.
[4] The adoption studies contain no tradeoff between abating pollution and increasing the efficiency of production, while in the R&D studies general innovation is held offstage (and is presumably constant), with the firm being represented by an abatement cost function which can be shifted by undertaking investment in pollution-augmenting

plementarity between productivity and abatement ignores the very real possibility that firms may face a tradeoff between pollution-saving innovation and general innovation. This so-called "crowding out" effect figures prominently in simulation studies of endogenous technical change (Goulder and Schneider, 1999; Popp, 2004a,b). Here we demonstrate rigorously that crowding out is pivotal to the influence of environmental regulation on the rate and direction of innovation, and provide insights into why.

The starting point for our analysis is Acemoglu's (2002) model of directed technical change, which represents a crucial breakthrough in the analysis of the tradeoffs among different lines of research. The central idea which motivated the early literature on ITC was the innovation possibility frontier, which represented supposedly "fundamental" tradeoffs in the augmentation of one factor of production relative to another. Acemoglu's seminal contribution is to elucidate how the tradeoff between different lines of research emerges endogenously out of a dynamic optimization framework, and depends on firm and market characteristics. This insight provides an unprecedented opportunity to formally elaborate Hicks's original intuition in the context of a regulatory constraint on the firm, and to relate the results to recent debates on the Porter Hypothesis and the crowding out.

Applications of Acemoglu's results to the environmental arena (e.g., Smulders and de Nooij, 2003) have retained his growth-theoretic framework: a two-sector economy with "clean" and "dirty" inputs, whose productivities each depend on the quantity and quality of a continuum of complementary intermediate goods ("machines"). The solution to this model illustrates that while a quantitative limit on the dirty input induces a pollution-saving bias of technical progress, the associated crowding out of R&D can actually reduce long-run growth. The present paper demonstrates how the same kind of result emerges from a simplified framework in which a firm's innovation responds to the changes in the relative prices of its inputs as a consequence of environmental regulation.

Our simple alternative to Acemoglu's abstract production structure is an intertemporally optimizing representative firm facing a downward-sloping demand curve for its output, which it produces using a clean and a dirty good. The firm augments each kind of input by investing in the appropriate kind of R&D. This model captures the essence of the Goulder and Schneider's and Popp's simulations while transparently elucidating the situation originally envisaged by Hicks. The desired tradeoff between clean- and dirty-augmenting R&D emerges out of the tension between the desire for profit-enhancing input augmentation and the need to undertake costly research to generate innovations, both of which adjust to

research. Popp (2005) shows that when the outcome of R&D is uncertain and follows a Pareto distribution, a tax on pollution which reduces a firm's expected profit will induce innovation which has a high probability of lowering profit even further but which nonetheless has a low probability of increasing profit above its pre-tax level.

changes in the relative price of inputs. The model is sufficiently tractable that all of its variables may be expressed as closed-form functions of a tax on pollution, which raises the relative price of the dirty input, alters the relative attractiveness of pollution-augmenting R&D, and shifts the firm's innovation possibilities.

The results elaborate and extend the key early findings of Magat (1976, 1978, 1979) to an intertemporal context, shedding light on how crowding out depends on the firm's characteristics, those of its market environment and the stringency of regulation, and illustrating the consequences for profit and emissions. When R&D is subject to diminishing returns, even though the tax may increase pollution-augmenting innovation, the crowding out of general innovation eliminates the possibility of a technological free lunch. This finding casts the shortcomings of the Porter Hypothesis into sharp relief. Finally, by comparing the solution to the behavior of an identical firm whose technology is held constant we are able to characterize how ITC can lower the cost of achieving a given reduction in pollution, and explain Smulders and de Nooij's counterintuitive result that with ITC regulation may end up reducing output in the long run relative to a situation where innovation is exogenous.

The rest of the paper is organized into four sections. Section 2 sets up the framework to be used in the subsequent analysis, and elucidates the conceptual linkages between early theorizing on ITC and recent debates over crowding out. Section 3 lays out the details of the formal model and its solution. The key results are presented and discussed in section 4. Section 5 concludes with a summary of the main points, and a discussion of caveats and possible extensions to the analysis.

2 Background and Motivation

Our approach is deliberately simple, and begins with the model outlined in Acemoglu (2003). A firm produces output Q using quantities X of two variable inputs, indexed by i: a clean good, C, and a dirty good, D ($i = \{C, D\}$). Input markets are assumed to be competitive, with C and D in perfectly elastic supply at prices p_C and p_D. The price of output is p. Production is of the constant elasticity of substitution (CES) variety, so that at each instant of time, t, the production function is given by:

$$Q(t) = \left[\omega_C\left(\alpha_C(t)X_C(t)\right)^{\frac{\sigma-1}{\sigma}} + \omega_D\left(\alpha_D(t)X_D(t)\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{1}$$

The parameters ω and σ denote the technical coefficients ($\sum_i \omega_i = 1$) and the elasticity of substitution between the input quantities measured in efficiency units. The variables α_i are augmentation coefficients which are under the firm's control, indicate the current state of input-augmenting technology, and are complementary to the input demands. Taking prices as exogenous, intertemporal

Fig. 1. The Innovation Possibility Frontier

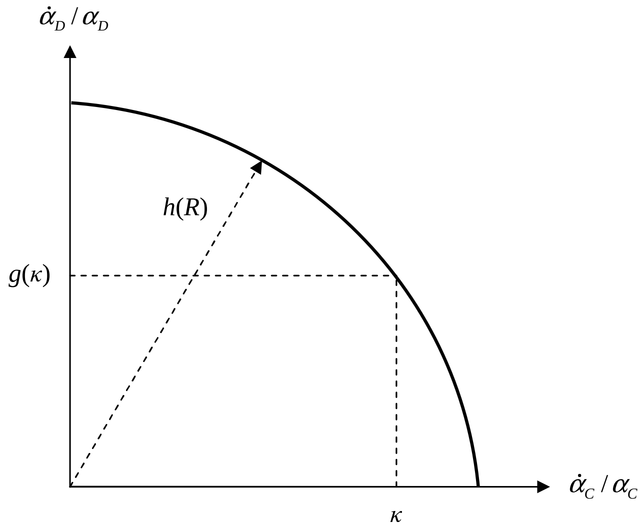

profit maximization over the planning horizon $t \in [0,\infty]$ implies that the firm solves the following problem:

$$\max_{X_C(t),X_D(t),\alpha_C(t),\alpha_D(t)} \int_0^\infty V(t)e^{-rt}dt,$$

subject to (1), where r is the firm's discount rate and

$$V(t) = p(t)Q(t) - p_C(t)X_C(t) - p_D(t)X_D(t). \tag{2}$$

is the firm's instantaneous variable profit.

This model does not possess an interior solution for the simple reason that the α_is are unbounded, i.e., the firm's innovation in either direction can be arbitrarily large. The crowding out phenomenon boils down to the nature of the constraints on the firm's ability to innovate, in particular the degree to which these factors render innovation costly to the firm and mandate C- and D-augmenting technical progress to be either substitutes or complements. In the early literature on ITC, this was accomplished by the conceptual device of the innovation possibility frontier (IPF), which is a reduced-form representation of "fundamental" tradeoffs in the firm's ability to exploit opportunities to improve the productivity of its different inputs (Ahmad, 1966; Kennedy, 1964; von Weizsacker, 1965).

The IPF is shown in Figure 1 by the heavy downward-sloping locus. Its curvature is defined by a function g, which, in the same way as the standard production possibilities frontier, reflects increasing opportunity costs of efforts to innovate more rapidly in a particular direction. [5] The IPF's position is determined by the firm's innovatory effort, for which instantaneous research spending, R, is a

[5] i.e., $g \geq 0$, and $g', g'' \leq 0$.

convenient proxy. R&D increases the distance of the IPF from the origin, subject to diminishing returns, which Kamien and Schwartz (1968) represent using the function h.[6] κ is a positive variable indicating the share of research allocated toward C-augmenting innovation, and is the key control variable used by the firm to set the direction of technical change:

$$\frac{\dot{\alpha}_D(t)/\alpha_D(t)}{\dot{\alpha}_C(t)/\alpha_C(t)} = \frac{g[\kappa(t)]}{\kappa(t)}.$$

The growth in the absolute magnitude of the augmentation of each input depends on the overall pace of technical progress determined by the firm's R&D, and is given by:

$$\dot{\alpha}_C(t)/\alpha_C(t) = \kappa(t)h[R(t)] \tag{3a}$$

$$\dot{\alpha}_D(t)/\alpha_D(t) = g[\kappa(t)]h[R(t)]. \tag{3b}$$

These two expressions form the basis for an augmented model, which is closed by incorporating the cost of conducting research, $\Phi[R]$, into the profit function. Instantaneous net profit is the difference between variable profit and research expenditures:

$$\pi = V - \Phi \tag{4}$$

and Φ may be thought of as the cost of adjusting a stock of knowledge capital whose reproduction is governed by the rate of investment in R&D.[7] We can therefore re-write the firm's problem as:

$$\max_{X_C(t),X_D(t),\kappa(t),R(t)} \int_0^\infty \pi(t)e^{-rt}dt,$$

subject to (1)-(4). This is essentially identical to the early model developed by Kamien and Schwartz (1968), and applied (with myopic expectations) in a regulatory context by Magat (1976).

The limitation of this model is of course the IPF itself. The function g is the crux of the problem, as it is completely heuristic in character, lacking rigorous microeconomic foundations while imposing an exogenous pattern of crowding out. Acemoglu's (2002) key insight was to demonstrate how g is the *outcome* of the firm's intertemporal profit maximization. Abstracting from the details, the main idea is to model R&D as heterogeneous, by splitting it into C-augmenting and D-augmenting research, R_C and R_D. Then, dropping time subscripts and recasting research spending as $\Phi = \Phi[R_C, R_D]$, the firm's problem becomes

$$\max_{X_C,X_D,R_C,R_D} \int_0^\infty \pi e^{-rt}dt \tag{5}$$

[6] i.e., $h, h' > 0$, $h'' \leq 0$.

[7] Following the adjustment cost literature, we assume that Φ is continuous, increasing and twice-differentiable: $\Phi, \Phi', \Phi'' > 0$.

subject to (1), (2), (4), and the analogue of (3):

$$\dot{\alpha}_i = h_i[R_i, \alpha_i], \tag{3'}$$

where the functions h_i reflect both the productivity of each line of research as well as the durability of the resulting inventions.

The solution to this new model implicitly defines the tradeoff between innovation possibilities, in the same way as the computational simulations of Goulder and Schneider and Popp. This alternative formulation makes clear that the competition between C- and D-augmenting technical change is a function of the relative contribution of each kind of innovation to the firm's output in (1), the relative cost of each type of R&D in (5), and the relative productivity and durability of the fruits of these investments in (3′). We go on to elaborate these relationships in the next section.

3 The Model

Our fundamental assumption that the intensity of input augmentation is determined by the state of technological knowledge within the firm.[8] Knowledge is a stock variable: it accumulates as a result of new blueprints or ideas created by research, but also decays over time due to obsolescence. Following eq. (3′), we assume that knowledge and research are both differentiated in character and input-specific, with C-augmenting R&D driving the accumulation of C-augmenting knowledge, and the same for the dirty input. Each augmentation coefficient is thus a stock of input-augmenting knowledge. We therefore model the functions h_i using a linear perpetual inventory formula:

$$\dot{\alpha}_i = \eta_i R_i - \delta \alpha_i, \tag{6}$$

in which the parameters δ and η_i reflect the decay of knowledge due to obsolescence and the productivity of each kind of R&D.[9] Eq. (6) is the core of Acemoglu's (2002) model with no state dependence. In the present context, the depreciation term implies that the firm must "run to stay in place", so that a steady state can only be achieved by continually investing in R&D.

[8] The firm's stock of knowledge can be thought of as the amalgam of the technical skills and managerial capabilities embodied in its workforce, technology embodied in its capital stock, and dismebodied patents, designs or codified organizational routines.

[9] Although spillovers are an important real-world aspect of ITC, we ignore them for the sake of analytical tractability and expositional clarity. For similar reasons we also assume that the depreciation rates of each kind of knowledge are the same. Relaxing either of these conditions causes the separability of R&D and innovation to break down, with the result that a closed-form solution of the model cannot be obtained.

One additional element is required to solve the problem in (5). It is not sufficient to solve the dual problem of minimizing the present discounted value of the unit cost of production—to compute the value of the augmentation coefficients we need to know the absolute quantities of each type of R&D, which depends on the size of the firm. The latter is determined by the equilibrium in the product market, which implies the need to explicitly model the demand for the firm's output. We do this by employing the commonly-used assumption (e.g., Baker and Shittu, 2006) of a downward-sloping demand curve for the firm's product, whose price elasticity is $\gamma > 1$:

$$Q = Mp^{-\gamma}. \tag{7}$$

The transformed problem is thus

$$\max_{Q, X_C, X_D, R_C, R_D} \int_0^\infty \pi e^{-rt} dt,$$

subject to (1), (2), (4), (6) and (7). [10]

The model is closed by assuming that research activities consume units of the final good. R&D exhibits increasing costs, which we model using a separable quadratic function (cf. Parry et al., 2003):

$$\Phi = \frac{1}{2} \sum_i \phi_i (1 + \psi_i) R_i^2. \tag{8}$$

The parameter ϕ_i reflects the costliness of i-augmenting research, while $-1 < \psi_i < 1$ is meant to capture either pre-existing taxes on either kind of research (as in Goulder and Schneider, 1999) or an exogenous R&D subsidy. Increasing research costs are consistent with diminishing returns, which previous theoretical and empirical studies have identified as a key aspect of the market for R&D (e.g., Jones, 1995; Popp, 2002).

Going back to the discussion in the introduction, we are careful to note that the combination of convex research costs in (8) and linear research productivity in (6) rules out the kinds of increasing returns which previous studies have found to be central to a Porter Hypothesis result. Nevertheless, the fact that both the cost and productivity of R&D are separable makes our model tractable enough to yield closed-form solutions for all of the firm's variables, which yields insights into the cost-savings associated with technological change.

The solution to the firm's problem is given in Appendix A. With the appropriate

[10] We acknowledge the inconsistency in modeling the firm's output price as endogenous while assuming that its input prices are exogenous, but we shall see that the gains from this sacrifice in terms of clarity and tractability immeasurably outweigh the costs of attempting to make the price of the dirty input (say) endogenous as well.

normalization of output the size of the firm is:

$$Q = \chi^{-\gamma}, \tag{9}$$

where χ is a CES unit cost function with input-augmenting technical change:

$$\chi = \left(\omega_C^\sigma \alpha_C^{\sigma-1} p_C^{1-\sigma} + \omega_D^\sigma \alpha_D^{\sigma-1} p_D^{1-\sigma}\right)^{\frac{1}{1-\sigma}}. \tag{10}$$

The unconditional input demands are:

$$X_i = \omega_i^\sigma \alpha_i^{\sigma-1} p_i^{-\sigma} \chi^{\sigma-\gamma} \tag{11}$$

and instantaneous variable profits are:

$$V = \chi^{1-\gamma}/(\gamma-1). \tag{12}$$

Finally, the control variables for the state of the firm's technology follow the nonlinear differential equation:

$$\dot{R}_i = (r+\delta)R_i - \frac{\eta_i \omega_i^\sigma \alpha_i^{\sigma-2} p_i^{1-\sigma} \chi^{\sigma-\gamma}}{(1+\psi_i)\phi_i} \tag{13}$$

Our results below are derived using the two stock evolution equations (6) in conjunction with the two costate evolution equations (13). To study ITC, we investigate the effects of an increase in the relative price of the dirty input due to the imposition of a tax τ on the dirty input. We employ the simplifying assumption that the units of ω_C and ω_D can be chosen to normalize pre-tax input prices to unity ($p_C = p_D = 1$). Then (slightly abusing notation) with environmental regulation, $p_D = \tau > 1$.

4 Results

4.1 Pollution taxes and the inducement of innovation

The dimensionality and nonlinearity of the differential equation system (6) and (13) complicate analysis of the transitional dynamics of the model. We defer such investigation to future research, and concentrate instead on the models' steady-state results (i.e., in which $\dot{R}_i = \dot{\alpha}_i = 0$), whose comparative statics are more transparent. In Appendix B we show that the second-order conditions imply that the ranges of the parameters σ and γ which are consistent with profit maximization are approximately $0 < \sigma < 2$ and $1 < \gamma < 2$.[11]

[11] We are not able to rigorously prove the stability of the steady-state for the corresponding ranges of σ and γ, owing to the difficulty of establishing whether the eigenvalues of

Using an asterisk ($*$) to indicate steady-state values, (6) and (13) reduce to:

$$\alpha_i^* = \eta_i R_i^* / \delta. \tag{14}$$

and

$$R_i^* = \frac{\eta_i \omega_i^\sigma (\alpha_i^*)^{\sigma-2} p_i^{1-\sigma} \chi^{\sigma-\gamma}}{(r+\delta)(1+\psi_i)\phi_i} \tag{15}$$

Together, these expressions yield the steady-state relative quantity of dirty R&D ($\varrho^* = R_D^*/R_C^*$):

$$\varrho^* = \left(\frac{\eta^{\sigma-1}\omega^\sigma \tau^{1-\sigma}}{\phi\psi}\right)^{\frac{1}{3-\sigma}}, \tag{16}$$

where $\eta = \eta_D/\eta_C$, and $\phi = \phi_D/\phi_C$ and $\psi = (1+\psi_D)/(1+\psi_C)$ denote the relative efficiency cost, and taxation or subsidization of D-augmenting R&D, and $\omega = \omega_D/\omega_C$ indicates the relative importance of the dirty input in production. Eq. (16) expresses the composition of R&D, and therefore the direction of innovation in the steady state, as a function of relative prices, formalizing Hicks's original intuition. It leads directly to the first result:

Proposition 1 *A tax on the dirty input induces a decrease (increase) in D's relative share of research when substitution among inputs is (in)elastic.*

In Table 1 we show precisely how τ's influence on the direction of technical progress depends on the value of the elasticity of substitution. The larger the value of σ the smaller the denominator of the exponents, leading to an amplification of the influence of prices on ϱ^*. The table also summarizes the effects of non-price factors on the direction of innovation in (16), and the manner in which they depend on σ. If σ is less than (greater than) unity, the greater the relative efficiency of D-augmenting research the smaller (larger) the relative quantity of this kind of R&D. Given the constraint on the feasible range of elasticity of substitution, the direction of the impacts of ϕ, ψ and ω are independent of the value of σ, with lower relative costs of, or larger taxes or smaller subsidies on, D-augmenting R&D, or a larger coefficient on the dirty input inducing a larger increase in D's share of total research.

The simple intuition behind the response of R&D to the tax is as follows. Because the firm's demand for the taxed good declines as the latter's price increases, research which augments this input generates a smaller increase in output and profit compared to R&D which augments the untaxed good. Thus, if D is not a necessary input to production (which is the case when $\sigma > 1$), the firm has an incentive to focus its research effort on augmenting the *un*taxed input, whose share of production expands with the rise in the taxed input's relative price. By contrast, when D is necessary ($\sigma < 1$), the firm behaves as predicted by Hicks's

the Jacobian of the dynamical system (6) and (13) have negative real parts. See Appendix C for details.

Table 1
Comparative Statics of Relative R&D Inducement in the Steady-State

	$\frac{\partial\varrho^*}{\partial\tau}$	$\frac{\partial\varrho^*}{\partial\eta}$	$\frac{\partial\varrho^*}{\partial\phi}$	$\frac{\partial\varrho^*}{\partial\psi}$	$\frac{\partial\varrho^*}{\partial\omega}$	$\frac{\partial\alpha^*}{\partial\eta}$
$0 < \sigma < 1$	+	−	−	−	+	+
$1 < \sigma < 2$	−	+	−	−	+	+

ϱ^* = Steady-state relative amount of D-augmenting R&D; α^* = Steady-state relative amount of D-augmenting innovation; τ = Tax on D; η = Relative productivity of D-augmenting R&D; ϕ = Relative cost of D-augmenting R&D; ψ = Relative taxation/subsidization of D-augmenting R&D; ω = Relative magnitude of the coefficient on D in production

conjecture, performing relatively more D-augmenting R&D to conserve on the input whose relative price is increasing.

Our findings put to rest a key debate in the early literature on ITC. Kennedy (1964) and Binswanger and Ruttan (1978) argued that innovation should be biased toward the input with the larger relative share of production, while Salter (1960) and Samuelson (1965) countered that the managers of the firm should be less interested in economizing on one input versus another than in reducing costs *in toto* (i.e., neutral technical progress). The beauty of eq. (16) is that it integrates both perspectives to provide the complete picture: the results are derived from an intertemporal optimization framework in which steady state relative quantity of i-augmenting R&D is increasing in the relative magnitude of the coefficient on input i.

Together, (14) and (16) pin down the steady-state relative quantity of dirty innovation ($\alpha^* = \alpha^*_D/\alpha^*_C$), in a manner comparable to Acemoglu (2002, eq. 26): [12]

$$\alpha^* = \left(\frac{\eta^2\omega^\sigma\tau^{1-\sigma}}{\phi\psi}\right)^{\frac{1}{3-\sigma}}. \tag{17}$$

This expression is the steady-state analogue of Kennedy's IPF, and elucidates how the degree of crowding out depends on the level of the pollution tax and the characteristics of the firm. While the effects of D's relative price, relative research cost, relative taxation and relative importance in production on the relative quantity of dirty innovation are identical to those in (16), Table 1 illustrates that the key difference is the effect of the relative efficiency of D-augmenting

[12] In Acemoglu's model the ratio of the rates of capital- and labor-augmenting innovation is a function of the relative magnitudes of the coefficients on labor and capital in production, the relative abundance of capital and labor, and the elasticity of substitution.

Table 2
Comparative Statics of Equilibrium Induced Bias With Respect to the Dirty Input

	$\frac{\partial\beta^*}{\partial\tau}$	$\frac{\partial\beta^*}{\partial\eta}$	$\frac{\partial\beta^*}{\partial\phi}$	$\frac{\partial\beta^*}{\partial\psi}$	$\frac{\partial\beta^*}{\partial\omega}$
$0<\sigma<1$	−	−	+	+	−
$1<\sigma<2$	−	+	−	−	+

β^* = Steady-state induced bias of technology w.r.t. D; τ = Tax on D; η = Relative productivity of D-augmenting R&D; ϕ = Relative cost of D-augmenting R&D; ψ = Relative taxation/subsidization of D-augmenting R&D; ω = Relative magnitude of the coefficient on D in production

R&D, whose influence on α^* is similar to the impact of D's importance on relative R&D.

We define the equilibrium induced bias of the firm's production technique as the shift in an input's relative share of production costs in the steady-state which is stimulated by ITC—i.e., endogenous changes in the values of α_C^* and α_D^*—holding input prices constant. From (11), the relative input demands and relative cost shares are:

$$X_D^*/X_C^* = \omega^\sigma(\alpha^*)^{\sigma-1}\tau^{-\sigma} \quad \text{and} \quad \tau X_D^*/X_C^* = \omega^\sigma(\alpha^*)^{\sigma-1}\tau^{1-\sigma}.$$

It is easy to see that, compared to the case where the values of α_D and α_C are fixed, the equilibrium induced bias with respect to D, β^*, is proportional to $(\alpha^*)^{\sigma-1}$:

$$\beta^* \propto \left(\frac{\eta^2\omega^\sigma\tau^{1-\sigma}}{\phi\psi}\right)^{\frac{\sigma-1}{3-\sigma}}. \tag{18}$$

This expression leads to the paper's second result, which was articulated by Magat (1979, Theorem 1):

Proposition 2 (Magat) *A tax on the dirty input induces a D-saving bias in the technique of production.*

The comparative statics of eq. (18) are summarized in Table 2. In a manner similar to Magat (1978), they demonstrate how ITC's influence depends on the elasticity of substitution. If D-augmenting research is relatively costly, then ITC lowers (raises) the dirty input's cost share if input substitution is (in)elastic. The influence of the relative productivity of D-augmenting research or the relative importance of D in production is symmetric and opposite.

While the foregoing results elucidate the effect of relative prices on the direction of innovation and the consequent induced bias of technology in the steady state, they are uninformative about the how regulation affects the firm's overall rate of technical progress. In particular, nothing so far implies that a rise in τ induces the firm to increase either the *absolute* quantity of D-augmenting R&D or the total quantity of research, as advocates of technology-forcing regulation suggest. Moreover, the fact that the environmental impact of the firm depends its use of the dirty input in absolute terms motivates us to compute the separate impacts of the tax on each line of research.

Using (14) and (15) in conjunction with (17), we derive closed-form expressions for the magnitudes of both augmentation coefficients in the steady state:

$$\alpha_C^* = k_1 \theta^{\frac{\sigma-\gamma}{(1-\sigma)(3-\gamma)}}. \tag{19a}$$

$$\alpha_D^* = k_2 \theta^{\frac{\sigma-\gamma}{(1-\sigma)(3-\gamma)}} \tau^{\frac{1-\sigma}{3-\sigma}}, \tag{19b}$$

where $\theta = 1+(\overline{\theta}-1)\tau^{\frac{2(1-\sigma)}{3-\sigma}}$, and $\overline{\theta}$, k_1 and k_2 are positive constants.[13] The levels of innovation depend only on firm and market parameters in addition to the price of the dirty input, and therefore allow the steady-state values of the marginal cost of production, output and profit to all be expressed as functions of τ. Moreover, with a stationary demand curve and stable input prices they are both constant, and the firm exhibits a constant level of output and profit, consistent with the existence of a steady state.

Proposition 3 *A tax on the dirty good has an ambiguous effect on the augmentation of the inputs to the firm. Its influence on the augmentation of the clean input is monotonic, with the sign of the effect depending on the relative magnitudes of the elasticity of substitution among inputs and the elasticity of demand for output. Its influence on the augmentation of the dirty input may be non-monotonic, and depends on the values of the elasticity of substitution, the elasticity of demand and the magnitude of the tax.*

Proposition 3 is an elaboration and extension of Magat's (1979) Theorems 3 and 4.[14] To shed light on its underpinnings, we consider the effect of relative prices on the absolute levels of augmentation by differentiating (19) with respect to τ:

$$\operatorname{sgn}\left[\frac{\partial \alpha_C^*}{\partial \tau}\right] = \operatorname{sgn}\left[\frac{\sigma-\gamma}{(3-\sigma)(3-\gamma)}\right] \quad \text{and} \quad \operatorname{sgn}\left[\frac{\partial \alpha_D^*}{\partial \tau}\right] = \operatorname{sgn}\left[\frac{1-\sigma}{3-\sigma} - \frac{\gamma-1}{3-\gamma}(\theta-1)\right].$$

As illustrated in Table 3, the restrictions on the feasible ranges of σ and γ in Appendix B imply that the first condition is negative once input substitution is in-

[13] $\overline{\theta} = 1+[\eta^2/(\phi\psi)]^{\frac{\sigma-1}{3-\sigma}} \omega^{\frac{2\sigma}{3-\sigma}}$, while $k_1 = \eta_C^{\frac{2}{3-\gamma}} [\delta(r+\delta)(1+\psi_C)\phi_C]^{\frac{-1}{3-\gamma}} \omega_C^{\frac{\sigma(1-\gamma)}{(1-\sigma)(3-\gamma)}}$ and $k_2 = [(1+\psi_C)\phi_C/\eta_C^2]^{\frac{\gamma-\sigma}{(3-\sigma)(3-\gamma)}} [\delta(r+\delta)(1+\psi_D)\phi_D/\eta_D^2]^{\frac{-1}{3-\gamma}} \omega_C^{\frac{2\sigma(\sigma-\gamma)}{(1-\sigma)(3-\sigma)(3-\gamma)}} \omega_D^{\frac{\sigma}{3-\sigma}}$.

[14] These state, respectively, that an effluent charge induces investment in pollution abatement innovation, with the likely (but not necessary) consequence being a reduction in output-enhancing innovation.

Table 3
Comparative Statics of Steady-State Input Augmentation and Innovation Cost

	$\frac{\partial \alpha_C^*}{\partial \tau}$	$\frac{\partial \alpha_D^*}{\partial \tau}$	$\frac{\partial \alpha_C^*}{\partial \psi_C}$	$\frac{\partial \alpha_C^*}{\partial \psi_D}$	$\frac{\partial \alpha_D^*}{\partial \psi_C}$	$\frac{\partial \alpha_D^*}{\partial \psi_D}$	$\frac{\partial \Phi^*}{\partial \tau}$
$0 < \sigma < 1$	$-$	?	$-$	$-$	$-$	$-$	$-$
$1 < \sigma < 2$	$\mathrm{sgn}[\sigma - \gamma]$	$-$	$-$	$\mathrm{sgn}[\sigma - \gamma]$	$\mathrm{sgn}[\sigma - \gamma]$	$-$	$-$

elastic ambiguous if it is elastic. In the first instance clean R&D falls in response to a tax on D, while in the second the direction of the effect depends on the sign of $\sigma - \gamma$. The situation is slightly more complicated in the case of the dirty input. If input substitution is elastic, an increase in τ causes D-augmenting innovation to decline. However, if substitution is inelastic the sign of the effect is ambiguous. In particular, an increase in the tax on dirty input induces an increase in the absolute quantity of D-augmenting innovation if $\tau < \tau_{\max}^{\alpha_D^*}$, where the threshold value of the tax is given by: [15]

$$\tau_{\max}^{\alpha_D^*} = \eta(\phi\psi)^{-\frac{1}{2}} \omega^{\frac{\sigma}{\sigma-1}} \left[\frac{(1-\sigma)(3-\gamma)}{(\gamma-1)(3-\sigma)} \right]^{\frac{3-\sigma}{2(1-\sigma)}} . \tag{20}$$

Above this threshold, an increase in τ causes α_D to decline.

The intuition behind this result is straightforward. If the inputs to production are highly fungible, then the firm has less incentive to innovate, as substitution is a relatively inexpensive margin of adjustment to the tax. Conversely, the more difficult it is to substitute C for D, the more the firm must rely on innovation to reduce costs, as the sign of $\partial \alpha^* / \partial \tau$ suggests. However, while the firm will always perform *relatively* more dirty R&D, whether or not its D-augmenting innovation increases in *absolute* terms depends on its cost structure. Eq. (20) can be thought of as a kind of Laffer curve: for $\tau < \tau_{\max}^{\alpha_D^*}$ the firm's costs of adjustment are sufficiently low that the increase in variable profits from innovation outweigh the costs of the additional research which it entails, while for $\tau > \tau_{\max}^{\alpha_D^*}$ the costs of adjustment are high enough that the firm finds it more profitable to slash research spending even as it allocates a disproportionate share of R&D toward D-augmenting innovation.

A key implication of the results is that, given sufficient information about the firm, a regulator can set an environmental tax to induce the maximal quantity of pollution-saving innovation ($\tau = \tau_{\max}^{\alpha_D^*}$). The hump-shaped response of α_D^* to the tax not only gibes with economic intuition, it is an important qualification to the claim that more stringent environmental regulation induces more environment-friendly research. An arbitrarily large increase in τ is not assured to bring forth

[15] Note that the term in square brackets is strictly positive for $\sigma < 1$, making this condition well defined.

additional D-augmenting research. Rather, past $\tau_{\max}^{\alpha_D^*}$ the regulatory burden on the firm becomes so onerous that R_D^* begins to fall, and further tax increases will accelerate its decline, eventually reducing it below its pre-tax baseline level. [16]

Eq. (19) also provides an insight into the effects of combining technology policy and environmental policy. Table 3 shows that while subsidizing the R&D associated with a particular input always increases the rate at which that input is augmented, doing so also has the potential to reduce the rate of augmentation of the other input. The influence on each input is symmetric, and only arises when the value of the elasticity of substitution exceeds that of the elasticity of output demand. For other combinations of the parameters, R&D subsidies have a positive effect on both types of augmentation. [17] Finally, it is useful to note that with inelastic substitution, $\partial\tau_{\max}^{\alpha_D^*}/\partial\psi < 0$, which implies that the threshold below which the tax induces the firm to conduct more D-augmenting innovation only increases if technology policies favor dirty research by lowering the cost of R_D^* relative to R_C^*.

4.2 *Implications for the Porter Hypothesis*

The central implication of the previous result is that the Porter Hypothesis cannot hold if input substitution is elastic, or if environmental regulation is too stringent—in these cases the firm's total quantity of innovation unambiguously declines. But it turns out that even if substitution is inelastic and pollution taxes are low, diminishing returns to R&D ensures that the increase in α_D^* never completely offsets the fall in α_C^*, or vice versa. This result is the corollary to Smulders and de Nooij's (2003) main finding, and also constitutes an important qualification to Magat's (1979) Theorem 2. [18] It is demonstrated by computing the firm's steady state research expenditure using (8), (14) and (19):

$$\Phi^* = k_3\theta^{\frac{(1-\gamma)(3-\sigma)}{(1-\sigma)(3-\gamma)}}, \tag{21}$$

[16] Somewhat ironically, when $1 < \sigma < 2$ it is possible for the augmentation of the *un*taxed (i.e., clean) input to increase with the tax. But such growth does not persist forever, as α_C^* is bounded above by $\lim_{\tau\to\infty}\alpha_C^* = [\delta(r+\delta)(1+\psi_C)\phi_C/\eta_C^2]^{\frac{-1}{3-\gamma}}\omega_C^{\frac{\sigma(1-\gamma)}{(1-\sigma)(3-\gamma)}}$. Clearly, lower taxes or higher subsidies targeted toward research in the clean input give α_C^* more headroom to grow.

[17] Another implication is that for firms with this combination of parameters, it is possible for the regulator to simulate D-augmenting innovation by taxing C-augmenting research. See, e.g., Otto et al. (2006).

[18] The latter states that an effluent charge has an ambiguous impact on overall research expenditure.

where k_3 is a positive constant.[19] The variable Φ is a measure of firm's overall innovatory effort, which responds to an increase in the tax on D according to:

$$\operatorname{sgn}\left[\frac{\partial \Phi^*}{\partial \tau}\right] = \operatorname{sgn}\left[\frac{1-\gamma}{3-\gamma}\right].$$

For γ in the range consistent with profit maximization, the effect of the tax on the firm's total R&D budget is always negative, in stark contrast to the rosy picture painted by the Porter Hypothesis:

Proposition 4 *When R&D is subject to diminishing returns, a tax on the dirty input unambiguously reduces the firm's steady-state level of innovatory effort, leading to an increase in its unit cost of production, and declines in its output and variable profit.*

The firm's production cost, output and profit depend on the direct effect of τ on the demand for input as well as its indirect effects on the augmentation coefficients. The joint impact of these factors is found by substituting (19) into (10) to derive the steady-state marginal cost of production:

$$\chi^* = k_4 \theta^{\frac{3-\sigma}{(1-\sigma)(3-\gamma)}}, \tag{22}$$

where k_4 is a positive constant.[20] The dependence of marginal cost on the tax is given by:

$$\operatorname{sgn}\left[\frac{\partial \chi^*}{\partial \tau}\right] = \operatorname{sgn}\left[\frac{2}{3-\gamma}\right],$$

which is always negative. Moreover, from (9) and (12) it is evident that

$$\operatorname{sgn}\left[\frac{\partial V^*}{\partial \tau}\right] = \operatorname{sgn}\left[\frac{\partial Q^*}{\partial \tau}\right] = \operatorname{sgn}\left[-\frac{\partial \chi^*}{\partial \tau}\right].$$

Interestingly, the tax simultaneously reduces both variable profits and R&D expenditure, which is the opposite of the situation envisaged by advocates of environmental regulation. At the optimum, the firm will reduce its innovatory effort so as to balance marginal benefit (the incremental savings on R&D spending) against marginal cost (the incremental loss in variable profits). A robust result is that this equilibrium never results in profit which is higher than the pre-tax level. Using (22) to expand (12) yields:

$$V^* = k_5 \theta^{\frac{(1-\gamma)(3-\sigma)}{(1-\sigma)(3-\gamma)}}, \tag{23}$$

[19] $k_3 = \frac{1}{2}\delta^2[\delta(r+\delta)]^{\frac{-2}{3-\gamma}}[(1+\psi_C)\phi_C/\eta_C^2]^{\frac{1-\gamma}{3-\gamma}}\omega_C^{\frac{2\sigma(1-\gamma)}{(1-\sigma)(3-\gamma)}}$.

[20] $k_4 = [\delta(r+\delta)(1+\psi_C)\phi_C/\eta_C^2]^{\frac{1}{3-\gamma}}\omega_C^{\frac{2\sigma}{(1-\sigma)(3-\gamma)}}$.

where k_5 is a positive constant.[21] Using this result in conjunction with (21), makes net profit in (4) simply:

$$\pi^* = k_6 \Phi^*, \tag{24}$$

where k_6 is a positive constant.[22] The implication is that the dependence of this expression on τ is identical to that of research expenditures, i.e., always negative.

These forgone profits are the cost of abating pollution. The steady-state unconditional demand for D as a result of the tax is:

$$X_D^* = k_7 \theta^{\frac{2(\sigma-\gamma)}{(1-\sigma)(3-\gamma)}} \tau^{-\frac{1+\sigma}{3-\sigma}}, \tag{25}$$

where k_7 is a positive constant.[23] The expression above is unambiguously declining in τ:

$$\text{sgn}\left[\frac{\partial X_D^*}{\partial \tau}\right] = \text{sgn}\left[-\frac{1+\sigma}{3-\sigma} - \frac{1+\gamma}{3-\gamma}(\theta - 1)\right].$$

We can therefore state our fifth result, which, given the neoclassical character of the firm should come as no surprise:

Proposition 5 *In general, a tax on the dirty input simultaneously reduces the firm's pollution and net profit. However, with inelastic input substitution it is possible to obtain higher profit and lower pollution by combining a sufficiently small tax with a sufficiently large subsidy to D-augmenting R&D.*

Although this result is foreordained by the structure of our model, it nonetheless demonstrates the point that environmental regulation which achieves a real reduction in pollution cannot increase a regulated firm's profit, even in the presence of endogenous technological change, without some associated positive externality or source of increasing returns. Furthermore, it suggests the absence of complementarity between innovation and abatement in the neoclassical model causes pollution and profits to always move in the same direction in response to a regulatory stimulus.

A natural question is whether a technology policy, when combined with environmental regulation, can generate the decoupling of pollution and profit envisioned by the Porter Hypothesis. The answer is yes, but only in very specific circumstances. To illustrate this result we examine the joint impact of the tax and a subsidy to D-augmenting R&D. The desired changes in profit and pollution can

---

[21] $k_5 = [\delta(r+\delta)(1+\psi_C)\phi_C/\eta_C^2]^{\frac{1-\gamma}{3-\gamma}} \omega_C^{\frac{2\sigma(1-\gamma)}{(1-\sigma)(3-\gamma)}} / (\gamma - 1)$.

[22] $k_6 = \frac{2r+\delta(3-\gamma)}{\delta(\gamma-1)}$, implying that profit is strictly positive if $1 < \gamma < 2r/\delta + 3$. Empirical studies of the obsolescence of knowledge estimate that $\delta \approx r$ (e.g., Pakes and Schankerman 1984), implying that $1 < \gamma < 5$, which is always true if the second-order conditions for profit maximization are satisfied.

[23] $k_7 = [(1+\psi_C)\phi_C/\eta_C^2]^{\frac{2(\sigma-\gamma)}{(3-\sigma)(3-\gamma)}} [(1+\psi_D)\phi_D/\eta_D^2]^{\frac{1-\sigma}{3-\sigma}} \omega_C^{\frac{4\sigma(\sigma-\gamma)}{(1-\sigma)(3-\sigma)(3-\gamma)}} \omega_D^{\frac{2\sigma}{3-\sigma}} [\delta(r+\delta)]^{\frac{1-\gamma}{3-\gamma}}$.

then be written in terms of the total derivatives of these variables with respect to τ and ψ_D:

$$d\pi^* = \frac{\partial \pi^*}{\partial \tau}(\tau - 1) + \frac{\partial \pi^*}{\partial \psi_D} d\psi_D > 0 \quad \text{and} \quad dX_D^* = \frac{\partial X_D^*}{\partial \tau}(\tau - 1) + \frac{\partial X_D^*}{\partial \psi_D} d\psi_D < 0.$$

Together, these imply the following condition on the necessary change in the subsidy as a function of τ:

$$-\frac{\partial \pi^* / \partial \tau}{\partial \pi^* / \partial \psi_D} < d\psi_D < -\frac{\partial X_D^* / \partial \tau}{\partial X_D^* / \partial \psi_D},$$

which is equivalent to $-2 < d\psi_D < -\Psi$, where

$$\Psi = \frac{(3-\gamma)(1+\sigma) + (\gamma+1)(3-\sigma)(\overline{\theta}-1)\tau^{\frac{2(1-\sigma)}{3-\sigma}}}{(\gamma-1)(3-\sigma) + (3-\gamma)(\sigma-1)(\overline{\theta}-1)\tau^{\frac{2(1-\sigma)}{3-\sigma}}}.$$

The condition above is not well posed for arbitrary values of the parameters—at some tax rates, $-\Psi < -2$ for feasible combinations of σ and γ. To rule out this possibility we assume that Ψ is positive and small, which constrains substitution to be inelastic and implies an upper bound on the tax:[24]

$$\tau < \left[\frac{(\gamma-1)(3-\sigma)}{(3-\gamma)(1-\sigma)(1-\overline{\theta})} \right]^{\frac{3-\sigma}{2(1-\sigma)}}.$$

This establishes the result in proposition 5. It may be possible to obtain similar results for a general R&D subsidy, but we do not explore this possibility here. Finally, it is worth noting that this finding simply reinforces our main point that profits and pollution will tend to move in the same direction. It merely makes the free lunch (i.e., the subsidy revenue gained by the firm) explicit.

4.3 The cost savings from ITC versus exogenous innovation

The results thus far represent some unpleasant technological arithmetic. Nevertheless, we emphasize that the impossibility of a technological free lunch is in no way cause for pessimism. On the contrary, simulation studies suggest that the ability of innovation to respond to relative prices can significantly lower costs of pollution abatement compared to situations in which technology is constant. This section quantifies the potential of ITC in this regard, and sheds light on the origins of these benefits.

[24] For the tax to be meaningful, i.e., $\tau > 1$, we also require $\overline{\theta} < \frac{2(\gamma-\sigma)}{(3-\gamma)(1-\sigma)}$, which places additional restrictions on the parameters η, ϕ, ψ and ω. We assume that these are satisfied.

Our strategy is to demonstrate that the tax induces the firm to undertake more abatement, and simultaneously enables it to forgo less profit, with ITC compared to when technology is exogenous. Henceforth, we use a bar over a variable to indicate its value in the pre-tax steady-state, and a tilde (~) to indicate its steady-state value with the tax. The former is the baseline from which the incidence and the environmental benefit of the tax can be expressed as the fractional reductions in profit and the demand for the dirty input. When innovation is price-responsive these quantities are computed easily using (24) and (25):

$$\widetilde{\pi}^* / \overline{\pi}^* = \Phi^* / \overline{\Phi}^* = \left[1/\overline{\theta} + (1 - 1/\overline{\theta}) \tau^{\frac{2(1-\sigma)}{3-\sigma}} \right]^{\frac{(1-\gamma)(3-\sigma)}{(1-\sigma)(3-\gamma)}}. \tag{26}$$

$$\widetilde{X}_D^* / \overline{X}_D^* = \left[1/\overline{\theta} + (1 - 1/\overline{\theta}) \tau^{\frac{2(1-\sigma)}{3-\sigma}} \right]^{\frac{2(\sigma-\gamma)}{(1-\sigma)(3-\gamma)}} \tau^{-\frac{1+\sigma}{3-\sigma}}. \tag{27}$$

To investigate how the firm behaves differently in the absence of ITC, we consider a situation in which the firm's technology is exogenously fixed at the baseline pre-tax level, which we indicate using a caret (∧) over a variable. In particular, we assume that the managers of the firm do not optimize over R&D, but instead continue to conduct research at pre-tax levels after the tax is imposed, which renders both augmentation coefficients invariant to the regulation. The augmentation coefficients in (19) and total R&D spending in (8) therefore remain constant with $\theta = \overline{\theta}$:

$$\hat{\Phi}^* = \overline{\Phi}^* = k_3 \overline{\theta}^{\frac{(1-\gamma)(3-\sigma)}{(1-\sigma)(3-\gamma)}}. \tag{21$'$}$$

Making the appropriate substitutions, unit costs and variable profits in (10) and (12) become:

$$\hat{\chi}^* = k_4 \left[1 + (\overline{\theta} - 1) \tau^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{22$'$}$$

$$\hat{V}^* = k_5 \overline{\theta}^{\frac{(1-\gamma)(\gamma-\sigma)}{(1-\sigma)(3-\gamma)}} \left[1 + (\overline{\theta} - 1) \tau^{1-\sigma} \right]^{\frac{1-\gamma}{1-\sigma}}. \tag{23$'$}$$

Net profit and the demand for the dirty input are:

$$\hat{\pi}^* = \left\{ (k_6 + 1) \left[1/\overline{\theta} + (1 - 1/\overline{\theta}) \tau^{1-\sigma} \right]^{\frac{1-\gamma}{1-\sigma}} - 1 \right\} \hat{\Phi}, \tag{24$'$}$$

$$\hat{X}_D^* = k_7 \overline{\theta}^{\frac{(1-\gamma)(\gamma-\sigma)}{(1-\sigma)(3-\gamma)}} \left[1 + (\overline{\theta} - 1) \tau^{1-\sigma} \right]^{\frac{\sigma-\gamma}{1-\sigma}} \tau^{-\sigma}. \tag{25$'$}$$

Finally, the tax induces reductions in profit and pollution of:

$$\hat{\pi}^* / \overline{\pi}^* = \frac{1}{k_6} \left\{ (k_6 + 1) \left[1/\overline{\theta} + (1 - 1/\overline{\theta}) \tau^{1-\sigma} \right]^{\frac{1-\gamma}{1-\sigma}} - 1 \right\}, \tag{26$'$}$$

$$\hat{X}_D^* / \overline{X}_D^* = \left[1/\overline{\theta} + (1 - 1/\overline{\theta}) \tau^{1-\sigma} \right]^{\frac{1-\gamma}{1-\sigma}} \tau^{-\sigma}. \tag{27$'$}$$

We are now in a position to elucidate the difference made by ITC to the firm's response to the tax. We proceed by characterizing the differences between the steady-state tax loss functions (26) and (26′), and the marginal abatement cost functions (27) and (27′). While the nonlinearity of these expressions makes rigorous algebraic exposition difficult, it is possible to characterize the implicit relationship between pollution abatement and forgone profits in each case by eliminating the tax as a parametric variable. To do this we invert (26) and (26′) and then substitute the resulting expressions for τ in eqs. (27) and (27′) to yield:

$$\tilde{X}_D^*/\overline{X}_D^* = (1-1/\overline{\theta})^{\frac{1+\sigma}{2(1-\sigma)}} (\tilde{\pi}^*/\overline{\pi}^*)^{\frac{2(\sigma-\gamma)}{(1-\gamma)(3-\sigma)}} \left[(\tilde{\pi}^*/\overline{\pi}^*)^{\frac{(1-\sigma)(3-\gamma)}{(1-\gamma)(3-\sigma)}} - 1/\overline{\theta} \right]^{\frac{1+\sigma}{2(\sigma-1)}}, \tag{28}$$

$$\hat{X}_D^*/\overline{X}_D^* = (1-1/\overline{\theta})^{\frac{\sigma}{1-\sigma}} \left(\frac{1+k_6\hat{\pi}^*/\overline{\pi}^*}{1+k_6} \right) \left[\left(\frac{1+k_6\hat{\pi}^*/\overline{\pi}^*}{1+k_6} \right)^{\frac{1-\sigma}{1-\gamma}} - 1/\overline{\theta} \right]^{\frac{\sigma}{\sigma-1}}. \tag{28′}$$

To further characterize the firm's response to the tax in each case, we numerically simulate $(\tilde{\pi}^* - \hat{\pi}^*)/\overline{\pi}^*$ and $(\tilde{X}_D^* - \hat{X}_D^*)/\overline{X}_D^*$ as functions of τ for different values of the parameters r, δ, η, ϕ, ω_C, σ and γ. The results are summarized in Table 4. The first two columns show the inducement effect of the tax on C- and D-augmenting innovation, the third and fourth columns show the effect of the tax on pollution with ITC and the fixed technology, and the fifth and sixth illustrate the corresponding effects on profits of the two firms. The numerical analysis investigates the response of these variables to a range of values of the tax. We consider three ensemble cases, each of which examines the sensitivity of key variables to different combinations of the elasticities of substitution and demand: a base case in which both inputs are equally important in production (A: $\omega_C = \omega_D = 0.5$), one where the dirty input is the less important of the two (B: $\omega_D = 0.2$), and one where it is more important (C: $\omega_D = 0.8$). As expected, the pattern of results is similar across the cases, with its amplitude diminished in B and exaggerated in C.

As indicated by Table 3, the effect of the tax on innovation depends on the combination of values of σ and γ in a way that is highly nonlinear. For γ close to unity, low values of σ are associated with increases in D-augmenting at the expense of C-augmenting innovation in response to the tax, while for high values of σ the outcome is reversed. With a high value of γ, a low value of σ is associated with a decline in both types of innovation, and with a high value of σ α_D declines, but the response of α_C is monotonic in the direction of the sign of $(\sigma - \gamma)$. With inelastic input substitution, the value of τ at which α_D^* is maximized (shown in column 7) varies inversely with the demand elasticity. For $\gamma < 1.5$ the threshold $\tau_{\max}^{\alpha_D^*}$ is high enough that even many-fold increases in the relative price of D are likely to induce more D-augmenting research. However, in no case is the increase in one kind of innovation sufficient to offset the decline in the other.

Examining the effect of the tax on the demand for the dirty input, the firm abates

Table 4. The Effect of τ on Steady-State Profit and Pollution with ITC and Fixed Technology

γ	σ	$\alpha^*_C/\overline{\alpha}^*_C$ (%)				$\alpha^*_D/\overline{\alpha}^*_D$ (%)				$\tilde{X}^*_D/\overline{X}^*_D$ (%)				$\hat{X}^*_D/\overline{X}^*_D$ (%)				$\tilde{\pi}^*/\overline{\pi}^*$ (%)				$\hat{\pi}^*/\overline{\pi}^*$ (%)				$\tau^{\alpha^*_D}_{\max}$	$\tau^{\Delta X_D}_{\min}$	$\tau^{\Delta\pi}_{\min}$	$\tau^{\Delta\pi}_0$
	$\tau =$	1.1	1.5	2	5	1.1	1.5	2	5	1.1	1.5	2	5	1.1	1.5	2	5	1.1	1.5	2	5	1.1	1.5	2	5				
A. Base case: $\omega_C = \omega_D = 0.5$																													
1.1	0.5	−1	−5	−9	−21	1	3	5	9	−8	−30	−45	−76	−7	−28	−44	−75	−1	−2	−4	−9	−1	−2	−4	−9	28.4	2.4	1.8	3.2
1.1	1.5	1	6	9	18	−2	−8	−14	−31	−12	−43	−63	−90	−12	−41	−60	−88	−1	−2	−3	−6	−1	−2	−3	−6	–	2.3	1.3	1.8
1.5	0.5	−3	−11	−18	−39	−1	−3	−6	−16	−10	−37	−56	−86	−9	−34	−52	−83	−3	−13	−22	−46	−3	−12	−20	−44	0.3	2.3	3.5	18.1
1.5	1.3	−1	−3	−5	−10	−2	−10	−16	−33	−13	−46	−65	−91	−13	−43	−62	−89	−3	−12	−20	−37	−3	−11	−17	−34	–	1.9	4.7	109.5
1.5	1.7	1	4	6	11	−4	−17	−27	−54	−16	−54	−73	−96	−14	−48	−67	−93	−3	−11	−17	−28	−3	−10	−16	−29	–	1.8	1.7	3.4
B. Smaller coefficient on the dirty input in production: $\omega_C = 0.8, \omega_D = 0.2$																													
1.1	0.5	−1	−4	−7	−17	1	4	7	15	−7	−28	−43	−73	−7	−26	−40	−71	0	−2	−3	−7	0	−2	−3	−7	112.6	2.6	1.5	2.4
1.1	1.5	0	1	1	2	−3	−12	−20	−41	−14	−49	−68	−93	−13	−45	−64	−91	0	0	0	−1	0	0	0	−1	–	1.9	1.2	1.4
1.5	0.5	−2	−8	−14	−32	0	0	−1	−6	−9	−34	−51	−82	−8	−30	−47	−79	−2	−10	−17	−38	−2	−9	−15	−36	1.1	2.5	4.0	20.7
1.5	1.3	0	−1	−1	−2	−2	−8	−12	−26	−12	−43	−62	−89	−12	−42	−60	−88	−1	−3	−4	−9	−1	−2	−4	−8	–	2.1	12.8	∞
1.5	1.7	0	0	0	0	−5	−20	−31	−58	−18	−57	−76	−96	−15	−50	−69	−94	0	−1	−1	−1	0	−1	−1	−1	–	1.7	1.5	2.6
C. Larger coefficient on the dirty input in production: $\omega_C = 0.2, \omega_D = 0.8$																													
1.1	0.5	−2	−7	−11	−25	0	1	2	4	−8	−31	−48	−79	−8	−30	−47	−78	−1	−3	−5	−11	−1	−3	−5	−11	7.0	2.2	2.3	6.6
1.1	1.5	3	11	20	50	−1	−3	−5	−12	−10	−37	−55	−85	−10	−37	−54	−84	−1	−4	−7	−14	−1	−4	−6	−14	–	3.6	3.2	8.9
1.5	0.5	−3	−13	−22	−45	−1	−6	−10	−25	−12	−41	−60	−89	−10	−38	−56	−86	−4	−16	−27	−53	−3	−14	−24	−50	0.1	2.1	3.2	15.6
1.5	1.3	−1	−6	−9	−20	−3	−12	−20	−39	−14	−48	−68	−93	−13	−45	−64	−91	−6	−21	−33	−60	−5	−19	−30	−57	–	1.9	2.8	14.7
1.5	1.7	2	8	15	36	−3	−13	−21	−43	−15	−49	−69	−94	−13	−46	−65	−91	−6	−23	−36	−63	−5	−20	−33	−61	–	1.9	2.4	8.6

Throughout, we assume that $r = \delta = 0.1$, $\eta_C = \eta_D = \phi_C = \phi_D = 1$ and $\psi_C = \psi_D = 0$.

pollution most vigorously when its inputs are most fungible. Over the range of parameters considered, a ten-percent rise in D's relative price stimulates a reduction in pollution of 7-18 percent, while a five-fold increase in the relative price induces the firm to abate between 71-96 percent of its emissions. In all cases the reduction in pollution in response to a given value of τ is larger with ITC than with fixed technology. An interesting feature of the results is that the difference in pollution abatement between the two technological regimes first increases and then decreases τ grows larger. We return to this point below.

Turning now to the implications for profits, the less fungible the inputs to production and the more elastic the demand for output, the larger the loss precipitated by the tax. [25] The reduction in profit is modest for low levels of the tax (less than a six percent drop in response to a ten-percent rise in D's relative price), but increases rapidly with τ and can grow large (> 60 percent) for high values of the coefficient on D and the elasticity of demand. By far the most startling result is that ITC is actually associated with a *larger* economic losses, contrary to the intuition that the ability to innovate will enable the firms to reap higher profits for any value of the tax. Although seemingly anomalous, this result is actually easy to account for, as we explain below.

Interpretation of these results is facilitated by Figure 2, which reproduces their salient features in stylized form. Panel A plots the normalized difference in the firm's pollution and profit with ITC as compared to exogenous technology. As in Goulder and Schneider (1999), with ITC pollution abatement responds more elastically to the tax. Compared to the fixed technology, the ability of the firm to adjust its R&D lowers the cost of abatement, enabling the firm to use less of the dirty input for all values of τ. The behavior of the steady-state tax loss is less intuitive. Relative to the situation where technology is fixed, ITC enables the firm to enjoy higher profits only if the tax is above the threshold value of $\tau_0^{\Delta\pi}$—for taxes below this level technology's ability to adjust actually puts the firm at a profit *dis*advantage in the steady state.

The explanation for these phenomena is given in Panel B, where the heavy dashed loci denote the responses of the firm with fixed technology and the heavy solid loci indicate those of the firm with ITC. Quadrant I illustrates that the price-responsiveness of technology increases the convexity of the firm's marginal abatement cost function while simultaneously shifting it toward the horizontal axis, drastically increasing the decline in pollution at low values of tax. However, as τ grows large the firm's demand for D goes to zero regardless of whether innovation is price-responsive, with the result that the difference $(\widetilde{X}_D^* - \hat{X}_D^*)/\overline{X}_D^*$

[25] This result makes intuitive sense. The less capable the firm is of substituting the cheaper clean input for the more expensive dirty input, the more the burden of the tax is passed on to consumers through increases in the product price. In this situation, the more elastic the demand curve for output, the greater the drop in output which results.

Fig. 2. Changes in Profit and Pollution with Fixed Technology and with ITC

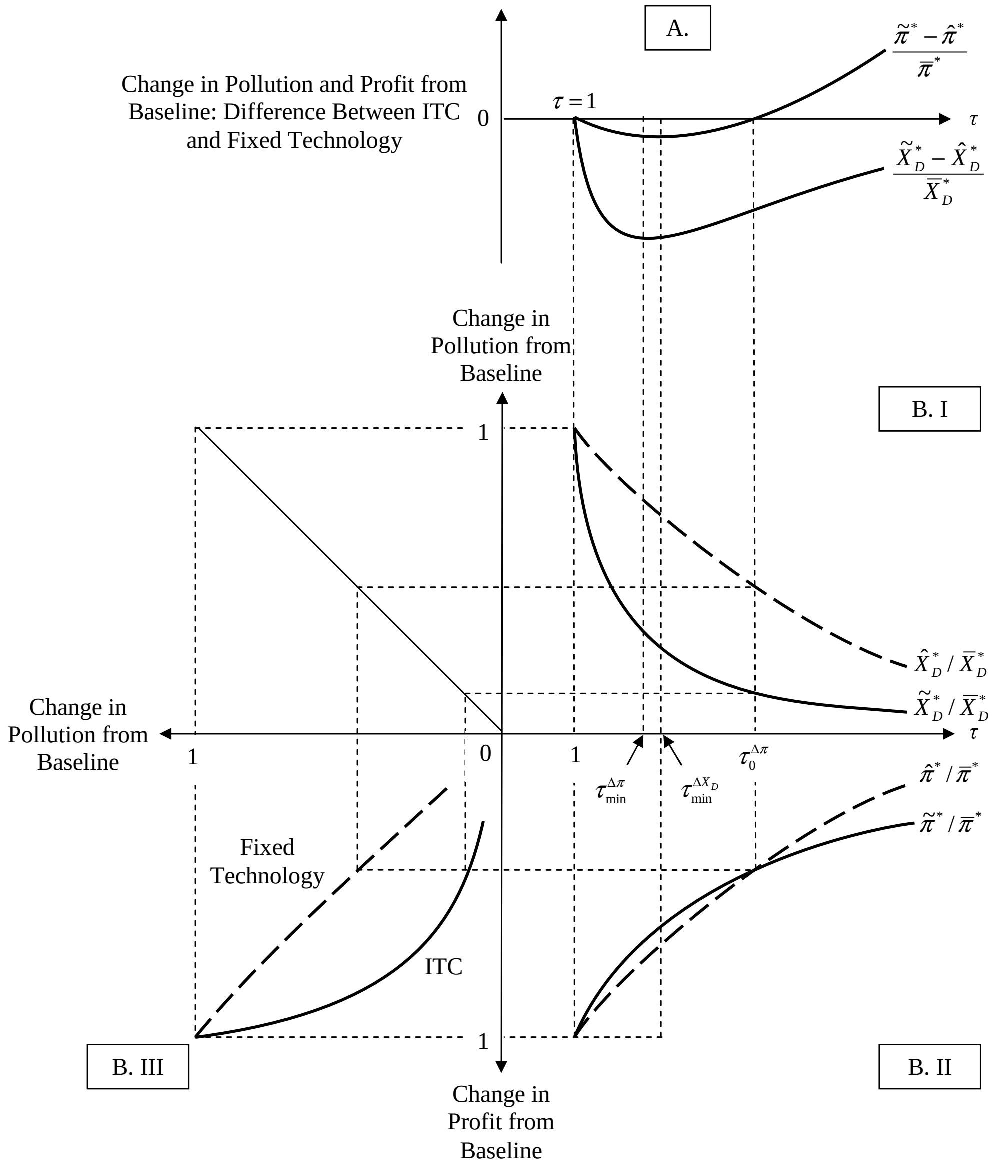

approaches zero asymptotically from below. Abatement with exogenous technology thus exhibits a minimum compared to that with ITC; at which point the value of the tax is given by $\tau_{\min}^{\Delta X_D}$. For the ranges of the parameters considered in Table 4, the nadir occurs at moderate levels of the tax (1.7-3.6). We summarize these results as follows:

Proposition 6 *In the steady-state, the reduction in pollution due to a tax on the dirty input is unambiguously larger with ITC compared to where innovation is exogenous.*

Quadrant II illustrates how the price-responsiveness of technology affects the profits forgone by the firm. ITC increases the convexity of the steady-state tax loss function, while simultaneously causing it to shift further away from the horizontal axis. Consequently, the firm's losses with ITC exceeds those with the fixed technology for $1 < \tau < \tau_0^{\Delta\pi}$. The reason for this behavior becomes clear if we examine the derivative of the tax loss function at the pre-tax baseline equilibrium:

$$\frac{\partial}{\partial\tau}\left(\frac{\widetilde{\pi}^* - \hat{\pi}^*}{\overline{\pi}^*}\right)\Bigg|_{\tau=1} = \left[\frac{\overline{V}^*}{\overline{\pi}^*}\frac{\partial}{\partial\tau}\left(\frac{\widetilde{V}^* - \hat{V}^*}{\overline{V}^*}\right)\Bigg|_{\tau=1}\right] - \left[\frac{\overline{\Phi}^*}{\overline{\pi}^*}\frac{\partial}{\partial\tau}\left(\frac{\widetilde{\Phi}^* - \hat{\Phi}^*}{\overline{\Phi}^*}\right)\Bigg|_{\tau=1}\right] \tag{29}$$

The second term in this expression is clearly negative by Proposition 4, while the sign of the first term is determined by

$$\mathrm{sgn}\left[\frac{\partial}{\partial\tau}\left(\frac{\widetilde{V}^* - \hat{V}^*}{\overline{V}^*}\right)\Bigg|_{\tau=1}\right] = \mathrm{sgn}\left[\frac{(\overline{\theta}-1)(\gamma-1)^2}{\overline{\theta}(\gamma-3)}\right]$$

which is also negative for γ in the appropriate range. Thus, recalling that positive baseline profits require that $\overline{V}^* > \overline{\Phi}^*$, eq. (29) says that with ITC the mitigating impact of research spending cuts is outweighed by the steeper decline in variable profit as productivity falls due to crowding out. The loss from an infinitesimally small tax is larger with ITC than in the case of fixed technology, a gap which (29) suggests initially increases with the tax.

We now examine the behavior of $(\tilde{\pi}^* - \hat{\pi}^*)/\overline{\pi}^*$ at the point where it crosses the horizontal axis in Panel A. To circumvent the nonlinearity of this expression we take a second-order Taylor series expansion of the difference between (26′) and (26) around $\tau = 1$, and set the result equal to zero. Solving for τ, the roots of the resulting equation are unity and, for values of σ and γ in the appropriate range, $\tau_0^{\Delta\pi} > 1$.[26] Table 4 shows that if input substitution is inelastic, or is elastic with $\sigma > \gamma$, then this point lies in the range 1.4-9. However, with elastic substitution and $\gamma > \sigma$ the value of the tax at which the steady-state deadweight losses are the same with ITC and with fixed technology is large (ranging from 14 to essentially infinite) and varies inversely with the importance of the dirty input in production. The implication here is that the difference between the steady-state deadweight losses with ITC and with fixed technology attains a minimum. The value of the tax where this occurs is indicated by $\tau_{\min}^{\Delta\pi}$, which in Table 4 lies in the moderate range 1.3-3.2 when input substitution is inelastic, or more elastic than demand. With elastic substitution and $\gamma > \sigma$, the nadir occurs at moderate to large values of the tax (between 2.8 and 12.8). Lastly, over the range of tax rates considered, the difference $(\widetilde{\pi} - \hat{\pi})$ is small, generally less than four percent of pre-tax steady-state profit.

These results account for Smulders and de Nooij's finding that a mandated reduction in the abundance of the dirty factor of production precipitates a decline

[26] The algebraic expression for the crossing point is a ratio of polynomials in σ, γ, $\overline{\theta}$ and k_6, and is too complicated to yield meaningful insights.

in the long-run growth of output with ITC relative to the situation where innovation is exogenous. It might seem counterintuitive, given that previous simulation results have always found lower deadweight losses with ITC. More importantly, the definition of a constrained optimum suggests that the present value of profits with ITC should always exceed that when technology is held constant. But in fact there is no contradiction, because the present result reflects a comparison *across steady states*. It is possible for the integral of the time-path of discounted profits with ITC to be larger, in spite of its lower long-run equilibrium level.

To see this, assume for simplicity that the firm's adjustment to the tax shock is such that over time its profit declines from the initial steady state value to its long-run equilibrium value at a constant rate, ξ, which is increasing in τ. Letting $\tilde{\xi}$ and $\hat{\xi}$ denote the values of this rate with ITC and with fixed technology, a larger present discounted value of profits with ITC requires:

$$\int_0^\infty \left[\left(\overline{\pi}^* - \hat{\pi}^*(\tau)\right) e^{-\hat{\xi}(\tau)t} + \hat{\pi}^*(\tau)\right] e^{-rt} dt \\ \leq \int_0^\infty \left[\left(\overline{\pi}^* - \widetilde{\pi}^*(\tau)\right) e^{-\widetilde{\xi}(\tau)t} + \widetilde{\pi}^*(\tau)\right] e^{-rt} dt,$$

which simplifies to the following expression for the upper bound on the steady-state profit differential:

$$\hat{\xi}(\tau)\left(\widetilde{\xi}(\tau) + r\right)\frac{\hat{\pi}^*}{\overline{\pi}^*} - \widetilde{\xi}(\tau)\left(\hat{\xi}(\tau) + r\right)\frac{\widetilde{\pi}^*}{\overline{\pi}^*} \leq \hat{\xi}(\tau) - \widetilde{\xi}(\tau).$$

This condition is satisfied if $0 < \widetilde{\xi} < \hat{\xi}\ \forall\ \tau > 1$. The implication is that, compared to the fixed technology, ITC enables the time-path of the firm's profit stream to decline relatively slowly in response to a given tax.

Only studies which characterize the firm's transition path are capable of capturing this outcome. Nevertheless, we caution that juxtaposing the present analytical results with those of computational simulations will not yield a clean comparison. In particular, the present model includes only variable inputs, not capital, while simulations tend to treat knowledge capital as a *substitute* for, rather than a complement to, pollution, and also employ a variety of heuristic devices to represent the tradeoffs between different kinds of research. [27]

[27] For example, in Popp's ENTICE model energy in efficiency units is a CES composite of carbon-energy and energy-saving knowledge capital, and in the Goulder and Schneider model knowledge capital substitutes for both energy and non-energy inputs at the top level of hierarchical CES production function. In ENTICE the opportunity costs and non-appropriability of energy-saving R&D result in a greater-than-proportional crowding out of physical capital investment, while Goulder and Schneider's model assumes that a carbon tax induces knowledge spillovers to producers of carbon-free energy, which mitigates the decline in the economy's overall rate of innovation.

In quadrant III we reflect the change in the demand for D onto the horizontal axis to obtain the firm's total abatement cost curves with ITC and the fixed technology. These portray the value of the tax loss at each level of emission reductions implicitly defined by (28) and (28′). Here the advantage of ITC is apparent—compared to the fixed technology, the steady-state tax loss is always smaller *for a given reduction in pollution*. Based on this result it is tempting to conclude that a firm facing an environmental tax smaller than $\tau_0^{\Delta\pi}$ will have an incentive to maintain its pre-tax pattern of innovation in order to earn higher steady-state profits, and will emit more compared to the situation in which it adjusted its research portfolio. The implication would then be that with endogenous technology a pollution standard would have lower costs than a tax, which runs counter to a large literature on the optimal choice of regulatory instruments when innovation is endogenous. However, the analysis above belies such reasoning—if the integral of discounted profits is larger with ITC then the firm will always choose to employ the additional technological margin of adjustment at its disposal. The results for the firm's profit may be summed up as follows:

Proposition 7 *In the steady state, the deadweight loss of the pollution tax with ITC is lower than that with exogenous technology only if the tax is above a certain threshold. Notwithstanding this, the forgone profit from reducing a given amount of pollution is always smaller with ITC than when technology is exogenous.*

5 Conclusion

This paper has analyzed the effect of environmental regulation on the rate and direction of technological change. We constructed a simple dynamic model of induced technical change in which a firm which faces a downward-sloping demand curve, produces output using clean and dirty inputs, and invests in clean and dirty research. A tax on pollution raises the relative price of the dirty input, increasing the relative attractiveness of pollution-augmenting R&D while crowding out clean R&D.

The results of the model allowed us to elaborate and extend key findings of the early literature on ITC. First, taxing the dirty input induces a decline in its relative share of research when input substitution elastic, and an increase in its share of research when substitution is inelastic. Second, the tax always biases the technique of production to conserve on the dirty good. Third, clean and dirty innovation are substitutes—the tax can induce more dirty innovation at the expense of clean innovation or vice versa. The former situation arises when the cheaper clean input is not readily substitutable for the more expensive dirty input. In this case the absolute quantity of dirty R&D exhibits a hump-shaped profile, increasing at first for small values of the tax and then declining thereafter.

We then went on to assess the implications for the Porter Hypothesis. Our fourth result is that even with the aforementioned inducement effect, a tax on the dirty input unambiguously reduces the firm's overall level of innovatory effort, leading to a rise in the cost of production and a fall in output and profit. Our fifth result follows as a logical consequence: the tax simultaneously reduces the firm's pollution and net profit. The exception to this rule is the special case where input substitution is inelastic and a sufficiently low pollution tax is coupled with a sufficiently large subsidy to D-augmenting R&D. These results turn on the fact that research is subject to diminishing returns, and does not enjoy either contemporaneous or intertemporal spillovers. They also demonstrate that the outcome envisaged by Porter Hypothesis cannot arise without some form of increasing returns or provision of free resources to the firm (e.g., subsidy revenue). Thus, introducing spillovers into the model is an important topic for further investigation.

Notwithstanding this finding, we demonstrated that ITC has the potential to lower the costs of environmental regulation. Our sixth result is that when innovation is price-responsive, the firm's abatement in response to a pollution tax is unambiguously larger than when innovation is exogenous. Seventh, however, in the steady state, the deadweight loss of the pollution tax with ITC is lower than that with exogenous technology only if the tax is above a certain threshold. Notwithstanding this, the forgone profit from reducing a given amount of pollution is always smaller with ITC than when technology is exogenous. These behaviors stem from the fact that the firm's response is more elastic with ITC, which we trace to the greater degree of convexity in the steady-state marginal abatement cost and tax loss functions. It is this phenomenon which accounts Smulders and de Nooij's key result that the reduction in the long-run growth rate of output is larger with ITC than with exogenous innovation. A key implication is that the time-path of the firm's response to a tax shock is flatter in the former case.

We close by noting that our conclusions are subject to a number of caveats. While we have focused on elucidating the attributes of the model's steady state and their implications, the transitional dynamics remain to be characterized. Moreover, the potential for increasing returns to decouple profits and pollution is an important issue that deserves further scrutiny. In this regard, the key challenge will be to find a way to incorporate spillovers while keeping the model tractable. More broadly, our focus on characterizing the mechanisms of ITC limits our analysis perhaps too narrowly to the firm. We consequently miss the feedback effects of changes in input demands on not only prices, but also the welfare impact of environmental damages, and the implications for cost-benefit analysis. Likewise, given our narrow consideration of a pollution tax, a natural extension of the current analysis is to consider the implications of ITC for the choice among instruments for environmental regulation. These are all topics which are ripe for investigation.

A The Solution to the Model

Using eq. (7) to substitute for p, the current-value Hamiltonian is:

$$\begin{aligned}\mathscr{H} = M^{\frac{1}{\gamma}} Q^{\frac{\gamma-1}{\gamma}} - p_C X_C - p_D X_D - \tfrac{1}{2}[(1+\psi_C)\phi_C R_C^2 + (1+\psi_D)\phi_D R_D^2] \\ + \lambda_C(\eta_C R_C - \delta\alpha_C) + \lambda_D(\eta_D R_D - \delta\alpha_D),\end{aligned}$$

where Q is given by (1), and λ_C and λ_D are the current-value adjoint variables dual to the knowledge stocks. Using the definition of the production function to make the appropriate substitutions, the first-order conditions are:

$$\frac{\partial\mathscr{H}}{\partial X_i}: \quad -p_i + \tfrac{\gamma-1}{\gamma} M^{\frac{1}{\gamma}} Q^{\frac{1}{\sigma}-\frac{1}{\gamma}} \omega_i \alpha_i^{\frac{\sigma-1}{\sigma}} X_i^{-\frac{1}{\sigma}} = 0$$

$$\Rightarrow \quad X_i = M^{\frac{\sigma}{\gamma}} \left(\tfrac{\gamma-1}{\gamma}\right)^{\sigma} Q^{\frac{\gamma-\sigma}{\gamma}} \omega_i^{\sigma} \alpha_i^{\sigma-1} p_i^{-\sigma} \tag{A.1}$$

$$\frac{\partial\mathscr{H}}{\partial R_i}: \quad -(1+\psi_i)\phi_i R_i + \eta_i\lambda_i = 0 \quad \Rightarrow \quad \lambda_i = (1+\psi_i)\phi_i R_i/\eta_i \tag{A.2}$$

$$\frac{\partial\mathscr{H}}{\partial \alpha_i}: \quad r\lambda_i - \dot{\lambda}_i = \tfrac{\gamma-1}{\gamma} M^{\frac{1}{\gamma}} Q^{\frac{1}{\sigma}-\frac{1}{\gamma}} \omega_i \alpha_i^{-\frac{1}{\sigma}} X_i^{\frac{\sigma-1}{\sigma}} - \delta\lambda_i \tag{A.3}$$

Eq. (A.1) represents the firm's optimal conditional input demands, which, when substituted into (1) yields the firm's level of output in terms of its unit cost of production:

$$Q = M\gamma^{-\gamma}(\gamma-1)^{\gamma}\chi^{-\gamma}. \tag{A.4}$$

where χ is given by eq. (10). To simplify notation we choose the units of output so that $M\gamma^{-\gamma}(\gamma-1)^{\gamma} = 1$. Substituting (A.4) into (A.1) yields the unconditional input demands in (11), which may be combined to yield the expression for variable profit in (12). Lastly, substituting (A.2) into eq. (A.3) along with (A.1) and (A.4), the equations of motion of the adjoint variables reduce to eq. (13).

B Second-Order Conditions for a Maximum

In this section we establish the ranges of key parameters which satisfy the second-order conditions for the existence of a maximum to the profit-maximization problem of the firm. From (A.1)-(A.3), the Jacobian vector of the Hamiltonian is:

$$\begin{bmatrix} \frac{\gamma-1}{\gamma}\omega_i \alpha_i^{\frac{\sigma-1}{\sigma}} X_i^{-\frac{1}{\sigma}} M^{\frac{1}{\gamma}} Q^{\frac{\gamma-\sigma}{\sigma\gamma}} - p_i \\ -(1+\psi_i)\phi_i R_i + \eta_i\lambda_i \\ \frac{\gamma-1}{\gamma}\omega_i X_i^{\frac{\sigma-1}{\sigma}} \alpha_i^{-\frac{1}{\sigma}} M^{\frac{1}{\gamma}} Q^{\frac{\gamma-\sigma}{\sigma\gamma}} - \delta\lambda_i \end{bmatrix}.$$

Using (1) to substitute for Q in this expression, we derive the Hessian, $\mathbf{H}^{\mathscr{H}}$, which is a 6×6 symmetric matrix whose lower-triangular entries are given by:

$$\begin{bmatrix}
\begin{array}{c} -z_0\omega_C\left(\alpha_C X_C\right)^{\frac{\sigma-1}{\sigma}} \\ \left\{[\gamma+\sigma(\sigma-1)]z_1+\sigma\gamma z_2\right\} \end{array} & & & & & \\
\begin{array}{c} z_0\omega_C\omega_D\alpha_C\alpha_D \\ (\gamma-\sigma)(\sigma-1) \end{array} & \begin{array}{c} -z_0\omega_C\left(\alpha_C X_C\right)^{\frac{\sigma-1}{\sigma}} \\ \left\{\sigma\gamma z_1+[\gamma+\sigma(\sigma-1)]z_2\right\} \end{array} & & & & \\
0 & 0 & -(1+\psi_C)\phi_C & & & \\
0 & 0 & 0 & -(1+\psi_D)\phi_D & & \\
\begin{array}{c} z_0\omega_C\left(\alpha_C X_C\right)^{-\frac{1}{\sigma}}(\sigma-1) \\ \left\{\sigma\gamma z_1+[\gamma+\sigma(\gamma-1)]z_2\right\} \end{array} & \begin{array}{c} z_0\omega_C\omega_D\alpha_D X_C \\ (\gamma-\sigma)(\sigma-1) \end{array} & 0 & 0 & \begin{array}{c} -z_0\omega_C\left(\alpha_C X_C\right)^{\frac{\sigma-1}{\sigma}} \\ \left\{[\gamma+\sigma(\sigma-1)]z_1+\sigma\gamma z_2\right\} \end{array} & \\
\begin{array}{c} z_0\omega_C\omega_D\alpha_C X_D \\ (\gamma-\sigma)(\sigma-1) \end{array} & \begin{array}{c} z_0\omega_D\left(\alpha_D X_D\right)^{-\frac{1}{\sigma}}(\sigma-1) \\ \left\{\sigma\gamma z_1+[\gamma+\sigma(\gamma-1)z_2]\right\} \end{array} & 0 & 0 & \begin{array}{c} z_0\omega_C\omega_D X_C X_D \\ (\gamma-\sigma)(\sigma-1) \end{array} & \begin{array}{c} -z_0\omega_D\left(\alpha_D X_D\right)^{\frac{\sigma-1}{\sigma}} \\ \left\{\sigma\gamma z_1+[\gamma+\sigma(\sigma-1)]z_2\right\} \end{array}
\end{bmatrix}$$

where $z_0 = \frac{\gamma-1}{\gamma\sigma^2} M^{\frac{1}{\gamma}} Q^{\frac{\gamma-\sigma}{\sigma}} / (z_1 + z_2) > 0$, $z_1 = \omega_C\,(\alpha_C X_C)\,(\alpha_D X_D)^{\frac{1}{\sigma}} > 0$, and $z_2 = \omega_D\,(\alpha_D X_D)\,(\alpha_C X_C)^{\frac{1}{\sigma}} > 0$. $\mathscr{H}$ attains a local maximum if the Hessian is negative semi-definite, a sufficient condition for which is that the k^{th}-order leading principal minors of $\mathbf{H}^{\mathscr{H}}$ have the same sign as $(-1)^k$. The leading principal minors are:

$$\begin{aligned}
m_1 =& -z_0\{z_2\gamma\sigma + z_1[\gamma+(\sigma-1)\sigma]\}X_C^{-\frac{\sigma+1}{\sigma}}\alpha_C^{\frac{\sigma-1}{\sigma}}\omega_C \\
m_2 =& z_0^2(z_1+z2)^2\gamma\sigma[\gamma+(\sigma-1)\sigma](X_C X_D)^{-\frac{\sigma+1}{\sigma}}(\alpha_C\alpha_D)^{\frac{\sigma-1}{\sigma}}\omega_C\omega_D \\
m_3 =& -(1+\psi_C)\phi_C m_2 \\
m_4 =& (1+\psi_C)(1+\psi_D)\phi_C\phi_D m_2 \\
m_5 =& z_0^3(z_1+z_2)^2\gamma^2\sigma^3\left\{z_2(\sigma-2)[\gamma+(\sigma-1)\sigma] + z_1[(\gamma-2)\sigma^2+2(\sigma-\gamma)]\right\} \\
& (\alpha_C X_C)^{-\frac{2}{\sigma}}(\alpha_D X_D)^{-\frac{\sigma+1}{\sigma}}(1+\psi_C)(1+\psi_D)\phi_C\phi_D\omega_C^2\omega_D \\
m_6 =& z_0^4(z_1+z_2)^4\gamma^3\sigma^5(\sigma-2)[(\gamma-2)\sigma^2+2(\sigma-\gamma)] \\
& (\alpha_C\alpha_D X_C X_D)^{-\frac{2}{\sigma}}(1+\psi_C)(1+\psi_D)\phi_C\phi_D\omega_C^2\omega_D^2.
\end{aligned}$$

Given the positivity of the share parameters, elasticities, unconditional input demands and output, m_1-m_4 have the appropriate signs if $\gamma + \sigma(\sigma-1) > 0$, while m_5 and m_6 have the appropriate signs if $\sigma - 2 < 0$ and $(\gamma-2)\sigma^2 + 2(\sigma-\gamma) < 0$. Together, these conditions imply that $\sigma < 2$ and $\sigma(1-\sigma) < \gamma < \frac{2\sigma(\sigma-1)}{\sigma^2-2}$, or the paired conditions $0 < \sigma < \sqrt{2}$, $\gamma > \frac{2\sigma(\sigma-1)}{\sigma^2-2}$ and $\sqrt{2} < \sigma < 2$, $\gamma < \frac{2\sigma(\sigma-1)}{\sigma^2-2}$. These automatically satisfy the assumption of elastic output demand, and imply that comparative statics analysis is valid over the range $0 < \sigma < 2$ and $1 < \gamma < 2$.

C Saddlepath Stability

The dynamical system (6) and (13) is stable if the eigenvalues of its Jacobian matrix all have negative real parts. The Jacobian is:

$$\mathbf{J} = \begin{pmatrix} -\delta & 0 & \eta_C & 0 \\ 0 & -\delta & 0 & \eta_D \\ j_1 & j_2 & r+\delta & 0 \\ j_3 & j_4 & 0 & r+\delta \end{pmatrix}$$

where

$$\begin{aligned}
j_1 &= \frac{\eta_C}{(1+\psi_C)\phi_C}\omega_C^\sigma \alpha_C^{\sigma-3}\chi^{2\sigma-\gamma-1}\left[\omega_C^\sigma \alpha_C^{\sigma-1}(2-\gamma)+\omega_D^\sigma \alpha_D^{\sigma-1}\tau^{1-\sigma}(2-\sigma)\right]\\
j_2 &= \frac{\eta_C}{(1+\psi_C)\phi_C}(\sigma-\gamma)(\omega_C\omega_D)^\sigma(\alpha_C\alpha_D)^{\sigma-2}\tau^{1-\sigma}\chi^{2\sigma-\gamma-1}\\
j_3 &= \frac{\eta_D}{(1+\psi_D)\phi_D}(\sigma-\gamma)(\omega_C\omega_D)^\sigma(\alpha_C\alpha_D)^{\sigma-2}\tau^{1-\sigma}\chi^{2\sigma-\gamma-1}\\
j_4 &= \frac{\eta_D}{(1+\psi_D)\phi_D}\omega_D^\sigma \alpha_D^{\sigma-3}\tau^{1-\sigma}\chi^{2\sigma-\gamma-1}\left[\omega_C^\sigma \alpha_C^{\sigma-1}(2-\sigma)+\omega_D^\sigma \alpha_D^{\sigma-1}\tau^{1-\sigma}(2-\gamma)\right].
\end{aligned}$$

Exploiting the fact that $r \approx \delta \ll 1$, so that $(r+2\delta)^2 \to 0$, the four eigenvalues of $\mathbf{J}$ may be written compactly as:

$$\frac{r}{2} \pm \frac{1}{2}\left\{2(\eta_C j_1+\eta_D j_4) \pm 2\left[(\eta_C j_1-\eta_D j_4)^2+4\eta_C\eta_D j_2 j_3\right]^{\frac{1}{2}}\right\}^{\frac{1}{2}},$$

the real parts of which cannot be signed for arbitrary values of the parameters.

References

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

Acemoglu, D., 2003. Factor prices and technical change: from induced innovations to recent debates. In: Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps. Princeton University Press, Princeton N.J., pp. 464–491.

Ahmad, S., 1966. On the theory of induced invention. Economic Journal 76, 344–357.

Ashford, N. A., 1994. An innovation-based strategy for the environment. In: Worst Things First? The Debate Over Risk-Based National Environmental Priorities. Resources for the Future, Washington D.C., pp. 275–314.

Ashford, N. A., Ayers, C., Stone, R. F., 1985. Using regulation to change the market for innovation. Harvard Environmental Law Review 9, 419–466.

Baker, E., Shittu, E., 2006. Profit-maximizing R&D in response to a random carbon tax. Resource and Energy Economics 28, 160–180.

Binswanger, H. P., Ruttan, V. W., 1978. Induced Innovation: Technology, Institutions, and Development. Johns Hopkins University Press, Baltimore.

Feichtinger, G., Hartl, R. F., Kort, P. M., Veliov, V. M., 2005. Environmental policy, the Porter Hypothesis and the composition of capital: Effects of learning and technological progress. Journal of Environmental Economics and Management 50, 434–446.

Goulder, L. H., Matthai, K., 2000. Optimal CO_2 abatement in the presence of induced technological change. Journal of Environmental Economics and Management 39, 1–38.

Goulder, L. H., Schneider, S. H., 1999. Induced technological change and the attractiveness of CO_2 abatement policies. Resource and Energy Economics 21 (3-4), 211–253.

Hicks, J., 1932. The Theory of Wages. Macmillan, London.

Jones, C. I., 1995. R&d-based models of economic growth. Journal of Political Economy 103, 759–784.

Kamien, M. I., Schwartz, N. L., 1968. Induced factor augmenting technical progress from a microeconomic viewpoint. Econometrica 37 (4), 668–684.

Kennedy, C., 1964. Induced bias in innovation and the theory of distribution. Economic Journal 74, 541–547.

Magat, W. A., 1976. Regulation and the rate and direction of induced technical change. Bell Journal of Economics 7 (2), 478–96.

Magat, W. A., 1978. Pollution control and technological advance: a dynamic model of the firm. Journal of Environmental Economics and Management 5 (1), 1–25.

Magat, W. A., 1979. The effects of environmental regulation on innovation. Law and Contemporary Problems 43 (1), 4–25.

Mohr, R. D., 2002. Technical change, external economies, and the Porter Hypothesis. Journal of Environmental Economics and Management 43 (1), 158–168.

Otto, V. M., Löschel, A., Reilly, J. M., 2006. Directed technical change and cost effectiveness of CO_2 reduction policies, paper presented at the Technological Change and the Environment Workshop, Dartmouth College, March 25-26.

Pakes, A., Schankerman, M., 1984. The rate of obsolescence of patents, research gestation lags and the private rate of return to research resources. In: R&D, Patents and Productivity. University of Chicago Press, Chicago, pp. 73–88.

Palmer, K., Oates, W. E., Portney, P. R., 1995. Tightening environmental standards: the benefit-cost or the no-cost paradigm? Journal of Economic Perspectives 9 (4), 119–132.

Parry, I. W., Pizer, W. A., Fischer, C., 2003. How large are the welfare gains from technological innovation induced by environmental policies? Journal of Regulatory Economics 54 (1), 39–60.

Popp, D. C., 2002. Induced innovation and energy prices. American Economic

Review 92 (1), 160–180.
Popp, D. C., 2004a. Entice: Endogenous technological change in the dice model of global warming. Journal of Environmental Economics and Management, 48(1) 48 (1), 742–768.
Popp, D. C., 2004b. R&D subsidies and climate policy: Is there a free lunch? NBER Working Paper No. 10880.
Popp, D. C., 2005. Uncertain R&D and the Porter Hypothesis. Contributions to Economic Analysis & Policy 4 (1), art. 6.
URL `http://www.bepress.com/bejeap/contributions/vol4/iss1/art6`
Porter, M. E., van der Linde, C., 1995. Toward a new conception of the environment-competitiveness relationship. Journal of Economic Perspectives 9 (4), 97–118.
Salter, W. E., 1960. Productivity and Technical Change. Cambridge University Press, Cambridge.
Samuelson, P. A., 1965. A theory of innovation along Kennedy-Weisacker lines. Review of Economics and Statistics 47 (4), 343–356.
Smulders, S., de Nooij, M., 2003. The impact of energy conversation on technology and economic growth. Resource and Energy Economics 25, 59–79.
Thirtle, C. G., Ruttan, V. W., 1987. The Role of Demand and Supply in the Generation and Diffusion of Technical Change. Fundamentals of Pure and Applied Economics. Harwood, Chur.
von Weizsacker, C.-C., 1965. Tentative notes on a two-sector model with induced technical progress. Review of Economic Studies 33, 245–251.
Xepapadeas, A., DeZeeuw, A., 1999. Environmental policy and competitiveness: the Porter Hypothesis and the composition of capital. Journal of Environmental Economics and Management 37, 165–182.