Paradigms and the Porter Hypothesis

Yoram Bauman*

January 28, 2004

1225 NE 61st St.

Seattle WA 98115

Tel: 206-351-5719

Email: yoram@smallparty.org

JEL codes: O31, O38, Q38

---

[*] Thanks to Ben Davis, Alan Haynie, Srobona Mitra, Karl Seeley, and Chad Westmacott for helpful comments.

Paradigms and the Porter Hypothesis

Abstract

In a 1995 debate on the Porter Hypothesis—the idea that tougher environmental standards can make companies more innovative and, ultimately, more profitable—Porter and van der Linde [13] argue that economists are locked into a "static mindset that environmentalism is inevitably costly." This paper analyzes the response from Palmer, Oates, and Portney [12] to show that Porter and van der Linde are, to a large extent, correct. We further argue that Palmer et al. provide strong (albeit anecdotal) evidence supporting the ideas in Thomas Kuhn's *Structure of Scientific Revolutions*, and discuss the importance of Kuhn's ideas for economics in general and the Porter Hypothesis in particular.

1 Introduction

The Porter Hypothesis asserts that stricter environmental standards can spur innovations that enhance competitiveness, and therefore that the right kinds of environmental policies can greatly reduce the costs of environmental policies and can even make companies more profitable. This paper provides an epilogue to a 1995 debate on the Porter Hypothesis that appeared in the *Journal of Economic Perspectives.*

As advanced by Porter and van der Linde [13], the Porter Hypothesis depends crucially on the dynamics of innovation:

> By stimulating innovation, strict environmental regulations can actually enhance competitiveness.... Fundamentally, [pollution] is a manifestation of economic waste.... [E]fforts to reduce pollution and [efforts to] maximize profits share the same basic principles, including the efficient use of inputs, substitution of less expensive materials and the minimization of unneeded activities.

Porter and van der Linde are quite specific in focusing on innovations in production practices, and contrast these with innovations in end-of-pipe waste treatment:

> [C]ompanies and regulators must learn to frame environmental improvements in terms of *resource productivity*, or the efficiency and effectiveness with which companies and their customers use resources. [Emphasis in original]

> . . . This view of pollution as unproductive resource utilization suggests a helpful analogy between environmental protection and product quality measured by defects. Companies used to promote quality by conducting careful inspections during the production process, and then by creating a service organization to correct the quality problems that turned up in the field. This approach has proven misguided. Instead, the most cost-effective way to improve quality is to build it into the entire process, which includes design, purchased components, process technology, shipping and handling techniques and so forth. This method dramatically reduces inspection, rework and the need for a large service organization. (It also leads to the oft-quoted phrase, "quality is free.")
>
> . . . *[Similar sorts of innovations in environmental quality are] central to our claim that environmental regulation can actually increase industrial competitiveness.* [Emphasis added]

In their conclusion, Porter and van der Linde accuse "economists as a group" of succumbing to the "static mindset that environmentalism is inevitably costly," and suggest that economists need to adopt a more dynamic perspective.

The response, defending economists and questioning the Porter Hypothesis, comes from Palmer, Oates, and Portney [12, hereafter referred to as POP]:

> Porter and van der Linde accuse mainstream environmental economics, with its "static mindset," of having neglected innovation. This charge is puzzling. For several decades now, environmental

> economists have made their case for incentive-based policy instruments (such as effluent charges or tradable emission permits) precisely by emphasizing the incentives that these measures provide for innovation in abatement technology.... Virtually every standard textbook in environmental economics makes the point that incentive-based approaches are perhaps more attractive for reasons of dynamic efficiency than for their ability to minimize the costs of attaining environmental standards at any particular point in time. A substantial literature has developed in recent years that explores the effects of various policy instruments on research and development decisions concerning abatement technology, a literature on which we shall draw in this discussion.

POP then provide a proof that tightening environmental regulations (which in their example means increasing the Pigovian tax rate) cannot benefit a profit-maximizing firm, regardless of opportunities for innovation.

This paper argues that Porter and van der Linde are, to a large extent, correct: environmental economists *have* by and large adopted a mindset that can appropriately be described as static in that it focuses on end-of-pipe abatement and overlooks opportunities for production process innovations. This "mindset" might be more accurately described as a paradigm: "a concrete piece of research or standard illustration that becomes a classic example of how 'good' science is conducted and that suggests further research" [1]. In this paper we take an in-depth look at one such standard illustration—the one Palmer et al. use to analyze innovation in pollution control—and argue that it provides strong (albeit anecdotal) evidence supporting the ideas in Thomas Kuhn's 1962 treatise, *The Structure of Scientific Revolutions*. We then discuss the importance of Kuhn's ideas for economics in general and the Porter Hypothesis in particular.

2 The Paradigm

The POP approach (shown in Figure 1 on page 6) features a firm currently using a technology with marginal abatement costs of MAC and assumes that "the firm could, if it chooses, reduce its marginal abatement cost function from the curve MAC to MAC^*" by spending some known amount on research and development (R&D).[1] The figure correctly identifies A and A^* as the optimal levels of abatement before and after innovation, respectively. To determine whether the R&D effort is worthwhile, POP assert that the firm must compare the benefits of the innovation with the (fixed) costs of R&D:

> [Figure 1] also depicts the gains to the polluting firm from undertaking the R&D effort, which can be divided into two parts. The source of the first part is that the earlier level of abatement activity becomes cheaper; the amount of gain here is given by triangle OFB.

---

[1] There is no uncertainty in this model. This issue—not the criticisms advanced in this paper—appears to be the focus of the caveat on page 122 of POP: "We emphasize that this model is static in character and fails to address the inherent uncertainty in [R&D] decisions. In this sense, it is subject to precisely the sort of criticism that Porter and van der Linde level in their paper."

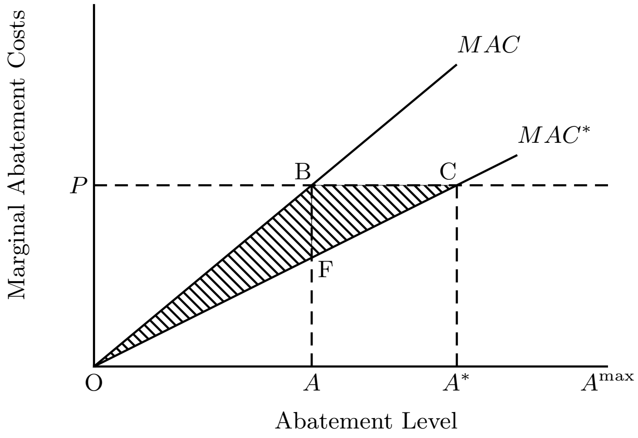

Figure 1: The standard graphical approach, with shaded area $OFCB$ representing the gain from innovation under a Pigovian tax of P.

> The second part comes from the new technology. The company will choose to abate a greater amount of pollution and thus avoid paying the pollution charge on that additional pollution; the gain here is the triangle BCF.
>
> The total gains to the polluting firm from innovation would thus be the area bounded by $OFCB$. Since the firm has not chosen this option, it must be that the cost of the R&D program that would move the firm from MAC to MAC^* exceeds the area of the profit that would be gained, $OFCB$.

This approach dominates the literature [2, 5, 11]. The following sections take a closer look at this approach's conclusion (that the area $OFCB$ measures the "gain" from innovation) and one of its assumptions (that innovations lower marginal abatement costs). We will show that this approach is “static” in that it is valid for innovations in end-of-pipe abatement efforts, but not for the production process innovations championed by Porter and van der Linde.

3 The Gain from Innovation

For expositional simplicity, we will assume that the innovation under consideration has no R&D costs or other fixed costs. In this case, the standard approach indicates that the area $OFCB$ measures the *net* gain from innovation, i.e., the profit differential between the pre-innovation firm and the post-innovation firm.

But there is an implicit assumption here, namely, that profits before and after innovation are equal at point O. This assumption can be seen from Figure 1 and the accompanying text, but it is also evident mathematically. The marginal benefit of abatement is P, the per-unit Pigovian tax, so the net impact of an additional unit of abatement (a) on the pre-innovation firm's profits (π) is $\frac{d}{da}\pi(a) = P - MAC(a)$. An identical result holds for the post-innovation firm's profits: $\frac{d}{da}\pi^*(a) = P - MAC^*(a)$. Making use of the Fundamental Theorem of Calculus, we can derive

$$\begin{aligned} OFCB &= OFCP - OBP && (1) \\ &= \int_O^{A^*} [P - MAC^*(a)]\, da \;-\; \int_O^{A} [P - MAC(a)]\, da && (2) \\ &= \int_O^{A^*} \left[\frac{d}{da}\pi^*(a)\right] da \;-\; \int_O^{A} \left[\frac{d}{da}\pi(a)\right] da && (3) \\ &= [\pi^*(A^*) - \pi^*(O)] - [\pi(A) - \pi(O)] && (4) \end{aligned}$$

$$= \quad [\pi^*(A^*) - \pi(A)] + [\pi(O) - \pi^*(O)] \tag{5}$$

This shows that the area $OFCB$ is equal to the profit differential $\pi^*(A^*) - \pi(A)$ if and only if $\pi(O) - \pi^*(O) = 0$, i.e., if and only if firm profits at point O are equal before and after innovation.[2]

Four comments are worth making here. First, point O is *not* the point of zero production. On the contrary, it is the point of zero abatement, and therefore can reasonably be thought of as the point of "full production": a firm facing no regulations would operate at point O.

Second, the term $\pi(O) - \pi^*(O)$ does *not* simply represent fixed costs. We have assumed fixed costs to be zero, but that does not imply that $\pi(O) - \pi^*(O) = 0$. If, for example, our costless innovation allows more electricity to be generated from the same quantity of coal and other inputs, profits after innovation will be higher at point O (and indeed at all other points at which the firm produces strictly positive amounts of electricity).

Third, profits at point O before and after innovation will be equal (i.e., $\pi(O) - \pi^*(O) = 0$) *if* the innovation is limited to end-of-pipe waste treatment technology. This is because point O represents the point of zero abatement, meaning that a firm producing at this level does not engage in any end-of-pipe waste treatment. Since there is no abatement activity and since we have assumed that there are no R&D costs, the profit levels at point O before and

---

[2] More generally, if the innovation has fixed costs of R, it can be shown that the area $OFCB$ is equal to the "gain from innovation" $[\pi^*(A^*)+R] - \pi(A)$ if and only if $\pi(O) - [\pi^*(O)+R] = 0$, i.e., if and only if firm profits pre- and post-innovation at point O differ only by the fixed costs R of the innovation.

after innovation must be equal.

Fourth and finally, profits at point O before and after innovation will (in general) be equal *only if* the innovation is limited to end-of-pipe waste treatment technology. If the innovation affects production processes (for example, if it allows more electricity to be generated from each ton of coal), there is every reason to think that profits at point O will be different: the firm may be using different inputs, or a different ratio of inputs, or may be squeezing more profits from its inputs. Indeed, the point O itself is likely to have moved.

Our analysis above shows that POP's approach is (in general) not valid for production process innovations. The next section identifies an additional difficulty in using the standard approach to critique the Porter Hypothesis: Porter-type innovations are unlikely to fit into the paradigm's assumptions about the effect of innovations on marginal abatement costs.

4 Innovation and Marginal Abatement Costs

One of the assumptions in the standard approach is that innovation lowers marginal abatement costs. For example, Downing and White [2] argue that

> *An innovation is a discovery that will reduce the costs of controlling emissions.* It normally involves an initial cost or investment (e.g., research and development expenses) and then a subsequent cost reduction or savings if the innovation is employed.... The cost savings from innovation can take several forms.... *The innovation which is most commonly discussed is one where both inframarginal*

> *and marginal units are less costly than the original cost function....*
>
> It is this last type of innovation that will be the focus of our analysis.
>
> [Emphases added.]

This type of innovation is also the focus of POP's analysis. But the innovations at the heart of Porter and van der Linde's claims might reasonably be expected to *increase* marginal abatement costs. This is because the marginal *cost* of *abating* pollution is also the marginal *benefit* of *emitting* pollution. Porter-type innovations enhance resource productivity, suggesting that innovating firms can squeeze greater value from each of their inputs, including the "input" of pollution.

To see this more clearly, we can rotate the marginal abatement cost curve in Figure 1 around the point of total abatement (or, more precisely, the line $a = A^{max}$) to get the marginal emissions benefits curve shown in Figure 2 on page 11. As with a demand curve, the area under a marginal emissions benefit curve between any two points (say, E^* and E) is the total benefit of increasing emissions from E^* to E. This total benefit is equivalent to the total cost of reducing emissions from E to E^*, i.e., increasing abatement from A to A^* in Figure 1.

Figure 2 provides some intuition to support the idea that production process innovations may increase marginal abatement costs. For a simple example, consider an innovation (motivated by Porter and van der Linde's product quality example) that allows a firm to generate more electricity for each ton of coal it uses. Such an innovation would *increase* marginal emissions benefits because

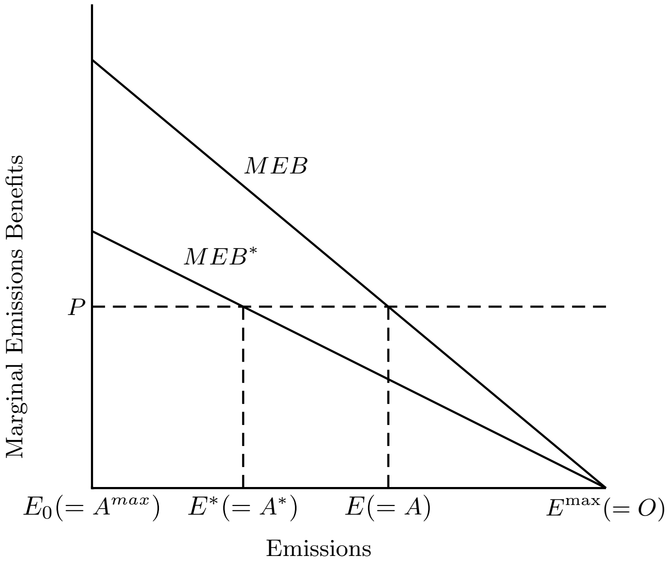

Figure 2: A marginal emissions benefit curve corresponding to the marginal abatement cost curve in Figure 1.

each unit of emissions provides the firm with a greater profit gain. Another example would be an innovation that allows the firm to use a less-polluting input (e.g., low-sulphur coal instead of high-sulphur coal). This innovation will "front-load" the demand for emissions, increasing the marginal benefits at low levels of emissions but decreasing the marginal benefits at higher level, as shown in Figure 3 on page 12.[3] Again, we see that innovation does not uniformly reduce

---

[3] For a numerical example corresponding to Figure 3, consider a firm switching to low-sulphur coal that has the same cost as high-sulphur coal but contains 50% less sulphur. The higher marginal benefits at lower levels of emissions arise because each unit of emissions corresponds to more activity; in particular, the first unit of emissions corresponds to twice as

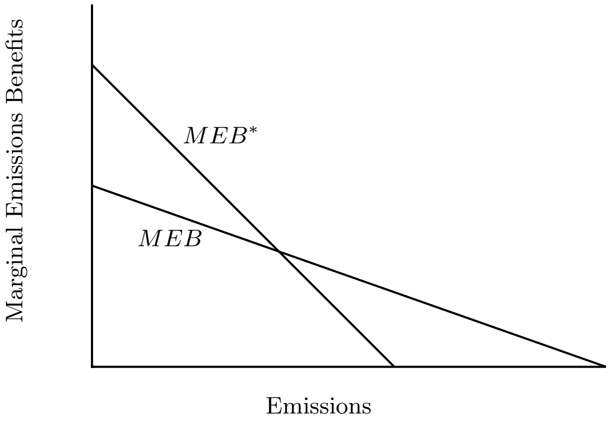

Figure 3: An innovation that "front-loads" demand for pollution.

marginal abatement costs.

Of course, innovation *will* uniformly reduce marginal abatement costs if the innovation concerns end-of-pipe waste treatment efforts. We again come to the conclusion that the standard illustration represents a "static mindset" that ignores production process innovations.

---

much coal as before, so the marginal benefits of the first unit of emission will double. The lower marginal benefits at higher levels of emissions arise because diminshing returns set in earlier *in terms of emissions*. As a final observation, note that, in the absence of regulation, firm profits should be identical before and after innovation: nothing changes except the coal's sulphur content, which is unpriced in the absence of regulation. This indicates that the areas under the two marginal benefit curves in Figure 3 should be equal. I am grateful to Karl Seeley for suggesting this graph and its interpretation.

5 Intermission

The previous sections show that POP's geometric approach lacks the generality necessary to address the production process innovations that are central to the Porter Hypothesis. Because of this, POP fail to prove their main contention: that tighter regulations, in the form of higher Pigovian tax rates, cannot benefit a profit-maximizing firm. Indeed, we hope to have turned the reader around 180° by showing that POP's paper *supports* Porter and van der Linde's contention about the "static mindset" of environmental economics.

We now hope to turn the reader around 180° again by showing that POP's claim—that higher Pigovian tax rates cannot benefit a profit-maximizing firm—is, in fact, true. To see this mathematically, let π^L and π^H denote the firm's maximum profits under "low" and "high" Pigovian tax rates, respectively; our task is to show that $\pi^L > \pi^H$. To do this, consider π^M, the profit level resulting from the *mimicry* situation in which the firm facing the low tax rate copies the behavior (i.e., the choices of inputs, outputs, *and technology*) of the firm facing the high tax rate. Note that $\pi^L \geq \pi^M$ because π^L is by definition the maximum profit for the firm facing the low tax rate. Note also that $\pi^M > \pi^H$ because the mimicking firm engages in the same behavior but faces a lower Pigovian tax rate; intuitively, the Pigovian tax rate plays the same role as an input price, and higher input prices hurt profits. Combining the two results produces $\pi^L > \pi^H$: profits are strictly greater with lower Pigovian taxes, regardless of opportunities for innovation.

Although $180° + 180° = 360°$, the world view the reader returns to may not be the same one she started out with. The concluding sections of this paper argue that what POP really demonstrate is the power of paradigms in economics—*and perhaps, by extension, in business.* Since paradigms lie at the heart of the Porter Hypothesis, we therefore hope to suggest to the reader that POP may have won the battle but lost the war.

6 An Economic Paradigm and its Evolution

What the debate over the Porter Hypothesis actually demonstrates is the power of the scientific paradigms first described in Thomas Kuhn's *The Structure of Scientific Revolutions.* Kuhn argued that scientists—far from being purely rational and unbiased—are in fact guided and constrained by paradigms that govern the process he called "normal science." The next section considers the possibility that businesses are similarly guided and constrained by paradigms of "normal business."

Kuhn's analysis does a remarkable job of explaining what might otherwise be interpreted as a neoclassical conspiracy against the environment. Regarding the resistance of environmental economists to the Porter Hypothesis and Porter and van der Linde's claims of a "static mindset," Kuhn writes (on pp. 5 and 7):

> Normal science, the activity in which most scientists inevitably spend almost all their time, is predicated on the assumption that the scientific community knows what the world is like. Much of the success of the enterprise derives from the community's willingness to defend that assumption, if necessary at considerable cost. Normal science,

> for example, often suppresses fundamental novelties because they are necessarily subversive of its basic commitments.... The commitments that govern normal science specify not only what sorts of entities the universe does contain, but also, by implication, those that it does not.

Elsewhere (p. 94), Kuhn helps explain POP's self-referential reliance on the (faulty) standard illustration of Figure 1:

> When paradigms enter, as they must, into a debate about paradigm choice, their role is necessarily circular. Each group uses its own paradigm in that paradigm's defense.

In the remainder of this section we describe the evolution of the paradigm used by POP. That history begins with early papers on innovation in pollution control—[2, 10, 16, 17]—all of which define innovation as a downward shift in marginal abatement costs or marginal emissions benefits. (The reference to "most commonly discussed" innovations in the Downing and White quote on page 9 suggests an even older history.) The use of the phrases "pollution control" and "pollution abatement" in these papers highlights an ambiguity that to some extent continues to this day: some authors interpret these as referring solely to end-of-pipe waste treatment costs, others to environmental innovation more broadly.[4]

---

[4] For example, Wenders [16, p. 393] focuses on "waste treatment costs", and so appears to use these terms specifically in relation to end-of-pipe innovations; but Zerbe [17, p. 371] uses an example of fuel switching, suggesting a broader scope. Over time, Zerbe's interpretation

Soon others adopt the same framework, in part because it is the existing framework: Malueg's (1989) defense of his use of downward-shifting marginal abatement cost curves includes the argument that "this is the formulation used by Downing and White. . . ." Over time, that framework becomes more and more deeply entrenched:

> Milliman and Prince [11]: "This paper expands the standard analysis of dynamic efficiency in pollution control. . . ."
>
> Jung, Krutilla, and Boyd [5]: "Milliman and Prince's study is the most comprehensive evaluation to date. . . . In this paper we extend [their] approach."
>
> Sunnevag [14]: "In examining the incentives for innovation with the use of voluntary instruments, I will follow the tradition of Downing and White (1986), Milliman and Prince (1989), and Jung et al. (1996)."

Concurrent with the gradual entrenchment of this framework is the gradual broadening of its scope. Downing and White [2, p. 19] explicitly define an innovation as "a discovery that will reduce the costs of controlling emissions" (p. 19) and take pains to point out some of the limitations of their approach; three years later, Milliman and Prince [11] use the same approach, minus most of the explicit definitions and caveats. Unwieldy phrases such as "technological

---

appears to have won out; see for example the textbook by Hanley, Shogren, and White [4, p. 108], in which the discussion of abatement cost functions considers changes in production processes as well as end-of-pipe waste treatment.

innovation in pollution control" are shortened to "technological innovation" or replaced with sweeter-sounding alternatives ("clean technology"), thereby helping to implicitly resolve—in favor of a broad interpretation—the ambiguities surrounding the terms "pollution control" and "pollution abatement."

Eventually the scope of the framework becomes all-encompassing:

> Hahn and Stavins [3], citing Milliman and Prince: "Incentive-based policies have been shown to be more effective in inducing technological innovation and diffusion…than conventional command-and-control approaches."
>
> Verhoef and Nijkamp [15], citing Milliman and Prince and others: "The general conclusion emerging from these studies is that 'economic' instruments (taxes or permits) usually provide larger incentives to adopt cleaner technologies than 'non-economic' instruments."
>
> Kim et al. [7], citing Milliman and Prince, Zerbe, Wenders, and others: "[T]ax-based regulatory policies provide the greatest incentive to search for and implement technological solutions, while the alternative of direct controls usually provides the smallest incentives to seek technological innovation."

This last paper is particularly noteworthy because its topic is nonpoint source pollution, where end-of-pipe abatement strategies are by definition impossible.

Having described an economic paradigm and its evolution, we now consider the implications for the Porter Hypothesis.

7 Paradigms in Business

POP turn to the standard approach of Figure 1 in their efforts to show that stricter regulations cannot make a profit-maximizing firm better off. Section 5 of this paper provides an alternative (algebraic) approach that verifies POP's assertion in a context that is more appropriate to the Porter Hypothesis debate than the geometric approach taken by POP. This alternative approach is also more concise and (at least arguably) easier, i.e., lower cost. One thing that it is not, however, is innovative: the approach is standard and the proof is straightforward. It would therefore be easy to say that POP could, if they chose, have adopted this approach.

Section 6's discussion of paradigms suggests that this view is misleading—that a more nuanced view is needed. This more nuanced view might also be profitably applied to the firm, about which POP write (on p. 123), "Let us now assume that the firm could, if it chooses, reduce its marginal abatement cost function. . . ." Such an assumption may not be warranted: the importance of paradigms in economics and other sciences indicates that paradigms may play a role in business decisions as well. Porter and van der Linde provide empirical support in their discussion of the quality revolution (quoted on page 2 above).

Our conclusion here is simply this: What separates Porter and van der Linde (and other such as Kemp [6]) from neoclassical economists is a belief about business paradigms. *To assume, as neoclassical economists do, that firms are profit-maximizing is to assume that firms have no paradigm at all, but rather an objective and unbiased vantage point from which to view the world and make* *decisions. This assumption should not be made lightly.*

Indeed, what POP have really "proved" (in the sense of "tested") is the power of paradigms—and the paradigms emerge victorious. Porter and van der Linde mount a direct and vivid assault on the established body of thought; they emphasize the importance of innovations in resource productivity; they contrast such innovations with end-of-pipe approaches; they assert that economists suffer under a "static mindset." In the face of this criticism, the response from POP is remarkable: "What distinguishes the Porter and van der Linde perspective from neoclassical environmental economics is *not* the "static mindset" of the latter. It is two other presumptions. . . " (emphasis added).

8 As in Science, So in Manufacture. . .

This paper makes four contributions to the Porter Hypothesis debate and economics more generally. The first is our demonstration that a "static mindset" does in fact distinguish much of neoclassical environmental economics from the Porter and van der Linde perspective.

Our second contribution is allegorical. We would suggest that assumptions about firms and their behavior make no more (and no less) sense than similar assumptions about economists. To paraphrase McCullough [9], "What is a firm, that an economist may know it, and an economist that he may know a firm?" This paper showcases the power of neoclassical assumptions such as profit maximization, and the concurrent limitations of those assumptions. As such, we believe that our story as a whole makes it seem more plausible that the private sector “systematically overlooks profitable opportunities for innovation.”

Our third contribution is Section 5’s algebraic alternative to the standard geometric approach. And there is an allegory here, too: our paper suggests that perhaps the analysis of totals (e.g., profits) should not be entirely subsumed by the marginal analysis that has become the focal point of economic teaching and practice. Margins are derivatives (both literally and figuratively) of totals, and this paper highlights the dangers of attempting to reverse that relationship.

Our paper’s final major contribution is to more clearly delineate the boundaries of the Porter Hypothesis debate. Our work does not come down strongly on one side of that debate or the other, but rather emphasizes the role and weight of profit-maximization or other assumptions about firm objectives and behavior. And the importance of these issues is not limited to the theoretical literature: Porter and van der Linde point out that many empirical papers on the Porter Hypothesis “cannot help reach the conclusions they do” because they begin by assuming that firms are profit-maximizing.

Forty years ago, Kuhn (p. 76) used an analogy to describe the resistance of “normal science” to paradigm shifts: “As in manufacture, so in science—retooling is an extravagance to be reserved for the occasion that demands it.” Now that Kuhn’s description of the scientific process is widely known and acknowledged, it may be time to turn that analogy around, yielding what amounts to a reformulation of the Porter Hypothesis: “As in science, so in manufacture—retooling is an extravagance to be reserved for the occasion that demands it.”

References

[1] George Argyrous. Kuhn's paradigms and neoclassical economics. *Economics and Philosophy*, 8:231–248, 1992.

[2] Paul B. Downing and Lawrence J. White. Innovation in pollution control. *Journal of Environmental Economics and Management*, 13:18–29, 1986.

[3] Robert W. Hahn and Robert N. Stavins. Incentive-based environmental regulation: a new era from an old idea. *Ecology Law Quarterly*, 18(1):1–42, 1991.

[4] Nick Hanley, Jason F. Shogren, and Ben White. *Environmental Economics in Theory and Practice*. Oxford Univ. Press, Oxford, 1997.

[5] Chulho Jung, Kerry Krutilla, and Roy Boyd. Incentives for advanced pollution abatement technology at the industry level: an evaluation of policy alternatives. *Journal of Environmental Economics and Management*, 30:95–111, 1996.

[6] Rene Kemp. *Environmental Policy and Technical Change*. Edward Elgar, Cheltenham, UK, 1997.

[7] C.S. Kim, N.D. Uri, C. Sandretto, and I.W. Parry. Endogenous technological change and pollution abatement. *Environmental Geology*, 27:191–197, 1996.

[8] Thomas S. Kuhn. *The Structure of Scientific Revolutions*. Univ. of Chicago Press, third edition, 1996.

[9] Warren S. McCulloch. *Embodiments of Mind.* MIT Press, Cambridge, 1989.

[10] Robert Mendelsohn. Endogenous technological change and environmental regulation. *Journal of Environmental Economics and Management*, 11:202–207, 1984.

[11] Scott R. Milliman and Raymond Prince. Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and Management*, 17:247–265, 1989.

[12] Karen Palmer, Wallace E. Oates, and Paul R. Portney. Tightening environmental standards: the benefit-cost or the no-cost paradigm? *Journal of Economic Perspectives*, 9(4):119–132, 1995.

[13] Michael E. Porter and Claas van der Linde. Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4):97–118, 1995.

[14] K. Sunnevag. Voluntary agreements and the incentives for innovation. *Environment and Planning C—Government and Policy*, 18:555–573, 2000.

[15] E.T. Verhoef and P. Nijkamp. Second-best energy policies for heterogeneous firms. *Energy Economics*, 21:111–134, 1999.

[16] John T. Wenders. Methods of pollution control and the rate of change in pollution abatement technology. *Water Resources Research*, 11:393–396, 1975.

[17] Richard O. Zerbe. Theoretical efficiency in pollution control. *Western Economic Journal*, 8:364–376, 1970.