Environmental Technology Policy and Emerging Technologies

WP 19.2
Energy Modeling Forum
Stanford University

Leon Clarke

April, 2002

Abstract

This paper explores the role and design of environmental technology policy when environmental innovation is embodied in an emerging environmental technology and when policy makers have been unable to fully internalize the environmental costs of production. The paper diverges from previous theoretical economic research on environmental innovation, first, by explicitly modeling the competition between two distinct technologies and, second, by using a simplified, general model of industry innovative behavior. The paper suggests that emerging environmental technologies fall into four possible categories, or "regions", each with its own welfare and policy implications. The paper also explores the allocation of research among technologies, and puts special emphasis on the distinct roles of demand-pull and technology-push policies.

1 INTRODUCTION

This paper is about the role and design of **environmental technology policy** when environmental innovation is embodied in **emerging environmental technologies**, or **EETs**. That is to say, it is about government policies to support the development and use of particular low-emissions technologies that are not currently competitive with their dominant competitors—technologies like photovoltaic cells, fuel cells, and electric cars. How well can we do using environmental technology policy? What sorts of environmental technology policies are best? Which technologies should governments support? Should governments rely primarily on demand-pull measures, such as technology standards and adoption subsidies, or on technology-push measures, such as R&D at government laboratories or subsidies for private sector R&D?

The research to follow falls within a set of theoretical economic literature that explores environmental policy in the presence of endogenous technological change using simple theoretical models[1]. Although this literature focuses primarily on the effectiveness of emissions instruments, such as emissions taxes, tradeable permits, and emissions quotas, a number of researchers have also explored technology policies. For example, Carraro and Soubeyran (1996) and Dosi and Moretto (1997) consider adoption subsidies, Carraro and Topa (1995) and Katsoulacos and Xepapadeas (1996) consider R&D subsidies, and Biglaiser and Horowitz (1995) consider technology standards[2].

The work in this paper will deviate from these previous efforts in three important ways. First, whereas previous research has relied on smooth parameterized production functions, upward sloping marginal abatement cost functions, and single, fixed-proportions technologies, I will specifically consider a situation of **technological competition**. Two distinct technologies, one polluting and one non-polluting, will compete to provide a generic good such as electricity or transportation. Second, I will use a general, simplified model of industry innovative behavior, the **industry planner model**, that parameterizes the factors—the **innovation market failures**, or **IMFs** for

[1]Notable works in this tradition include Magat (1978), Downing and White (1986), Milliman and Prince (1989), Jung et al. (1996), and Fischer et al. (1998).

[2]In an applied context, Goulder and Schneider (1997) and Goulder and Schneider (1999), and Duke and Kammen (1999) look at adoption subsidies. Newell (1997) and Newell et al. (1998) use econometric methods to capture the effects of past environmental policy on appliance efficiency.

short—that push private markets to invest in innovation in a socially nonoptimal fashion, regardless of environmental externalities[3]. Important IMFs are spillovers, deviations between private and social preference for risk, and deviations between private and social time preference. Previous research, in contrast, has assumed specific combinations of market structure, appropriability, research redundancy, and so forth. These assumptions both weaken the generality of the results and complicate analysis. Previous research has also neglected important issues of risk and time preference that are clearly relevant to environmental technology policy[4]. The industry planner model is general enough to implicitly include any IMF that causes the private sector to invest in innovation in a socially suboptimal fashion. The third difference between this and previous work is that I will explicitly focus on a situation in which the environmental costs of production have not been fully internalized by government policy—I will be exploring environmental technology policy under different variations of second-best.

This paper has two primary goals. The first goal is to increase our understanding of the role and design of environmental technology policy. Along these lines, I will show that there are four distinct possible scenarios, or "regions", in which EETs might find themselves, each with its own policy ramifications; I will emphasize the different roles of demand-pull policies and technology-push policies in each of these regions; and I will show that a policy program that begins with technology-push and ends with demand-pull is optimal in every region. The second goal is increase our understanding of the important characteristics of, and issues related to, the development and emergence of EETs. Along these lines, I will emphasize how the allocation of research effort among technologies, and not just the total level, is a relevant policy concern. I will show how incomplete internalization and IMFs bias innovation away from EETs and toward technologies currently in use.

The remainder of the paper is organized as follows. Section 2 will introduce the industry planner model. Sections 3 and 4 will explore the theoretical

[3]The industry planner model is consistent with an approach suggested by Nordhaus (2002).

[4]For example, a 1999 report on energy R&D from the President's Committee of Advisors on Science and Technology cited, as reasons that private sector R&D activities may fail to serve society's best interests, "... problems of market structure, magnitude of risk, size of investment, long time horizon for returns, or nonappropriability of research results by firms doing the research" (from PCAST (1999), p. 1-25). More is at issue than spillovers.

model of technological competition. Section 3 will look at a model in which innovation can only take place in the EET and will introduce the fundamental results of the analysis. Section 4 will extend the model to allow for innovation in both the EET and the dominant technology, and will focus on the importance of research allocation. Section 5 will sum up.

2 THE INDUSTRY PLANNER MODEL

In this section, I will present and discuss the industry planner model, a simplified, general model of industry innovative behavior. The model is a cornerstone of the analysis of EETs and technological competition in the following sections[5].

Two important questions in the analysis of technological advance are (1) what is the socially optimal level of innovation and (2) how will private industry invest in innovation? The industry planner model answers the second question by relating it to the first and explicitly breaking out the effects of incomplete internalization and IMFs. Industry innovation decisions are be made by a "flawed" planner, the industry planner, that serves as a proxy for industry innovative behavior.

The remainder of this section will proceed as follows. In Section 2.1, I will introduce the concept of **internal welfare**, an essential building block for the industry planner model. In Section 2.2, I will discuss socially optimal innovation and introduce the industry planner model. In Section 2.3, I will present the welfare metric, net welfare, I will use in the coming sections. Last, in Section 2.4, I will present extensions to, and generalizations of, the industry planner model.

2.1 Internal Welfare

This section will introduce the concept of internal welfare. Internal welfare is important because it is the basis for the welfare metric that the industry planner attempts to maximize through innovation.

Imagine an industry that produces a good, say widgets, and some concomitant emissions. Widgets might represent anything from electricity to

[5]The industry planner model is consistent with an approximate approach suggested by Nordhaus (2002).

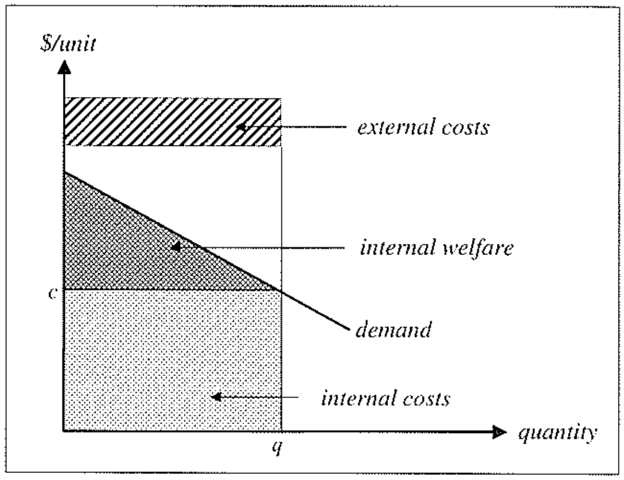

Figure 1: Social and Internal Welfare

domestic hot water to vehicle miles traveled. The market for widgets is competitive, so that prices are set at the private marginal production costs.

Let W be social welfare at a given production level—the social benefits of production net of the full social costs of production. Let w be internal welfare. Internal welfare differs from social welfare in that it is based only on internal production costs; it neglects external environmental costs. It follows that social welfare and internal welfare are linked by the following equality,

$$W = w - \text{external production costs} \tag{1}$$

Hence, if environmental costs have been internalized, internal welfare and social welfare will be identical. The divergence between the two welfare measures comes from incomplete internalization[6].

Figure 1 illustrates social and internal welfare when production is constant returns to scale. Let c be the private marginal costs of production. Internal welfare is total surplus as defined by these private costs: it is the area above the private cost curve and below the demand curve. Social welfare is this total surplus net of any uninternalized, or external, production costs, which are shown by the large block above the demand curve.

[6]I have defined internal and social welfare here assuming no distortions to demand. In reality, many policy measures, notably adoption subsidies, will shift the demand curve. In this case, internal welfare will be based on the *shifted* demand function and social welfare will be based on the *real* demand function. The equality in Equation 1 will not hold.

2.2 The Industry Planner Model and Socially Optimal Innovation

Now imagine that widget production technology might be improved through innovation. It is not important for now whether innovation alters mainly external environmental costs or internal costs; it only matters that there is potential for change. What is the socially optimal level of innovation and how will the market invest in innovation? This section provides the framework for answering these questions.

Assume that innovation takes place through conscious investments unrelated to production, loosely interpreted as R&D or design improvements[7]. Let ψ be the productive investment in innovation—the effort that really goes to improve widgets or widget production. Let internal and social welfare without innovation be given by w_0 and W_0, where the subscript "0" signifies "no-innovation". Let welfare associated with innovation level ψ be given by w_ψ and W_ψ. Again, if internalization is incomplete, social and internal welfare will differ.

The social planner wishes to find the level of innovation that will maximize the increase in social welfare net of innovation costs. Mathematically, the social planner solves the problem,

$$\max_{\psi} \quad NB_\psi = \overbrace{(W_\psi - W_0)}^{GB_\psi} - \psi \tag{2}$$

where NB_ψ represents the net social benefits of innovation level ψ, and GB_ψ represents the gross social benefits.

What level of innovation will private industry choose? I capture the behavior of private industry through a proxy planner, the industry planner, who solves the problem

$$\max_{\psi} \quad nb_\psi = \overbrace{\alpha(w_\psi - w_0)}^{gb_\psi} - \psi \tag{3}$$

[7]The important point here is that there are no explicit experience effects. Nonetheless, the research program could also be interpreted in the context of experience effects—research would be a production program aimed at both immediate revenues and at improving widget production.

where nb_ψ represents net industry planner benefits of innovation level ψ, and gb_ψ represents the gross industry planner benefits. α is a proxy parameter for IMFs. The lower is α, the more severe are the IMFs.

The idea in Equation 3 is that private innovative behavior differs from socially optimal innovative behavior through two forces, incomplete internalization and IMFs, and the effects are separable. Even if there were no IMFs, if $\alpha = 1$, private behavior might differ from socially optimal behavior because of incomplete internalization. The private sector would choose the level of innovation that optimizes the change in *internal welfare*, not the change in *social welfare*. Conversely, even if internalization were complete, if $w = W$, the private sector might still underinvest due to IMFs. This is captured by the parameter α. Figure 2 is a visual representation of these concepts.

The industry planner model is an attempt to model compactly three notions about private sector innovation. The first notion is that the private sector does not consider external costs in making innovation decisions. This notion is captured by having the industry planner optimize on internal rather than social welfare. The second notion is that private sector innovative behavior is fully positively correlated to the internal gains from innovation. That is to say, *ceteris paribus*, the higher are the internal gains, the greater will be the gains to private firms, and the greater will be the private incentive to innovate. The third notion is that private markets will underinvest in innovation in relation to the internal gains. In other words, even if the there were no external costs, the private sector would still not choose the optimal level of innovation because of IMFs. These last two notions are captured by making industry planner gross benefits a fraction, α, of internal gross benefits.

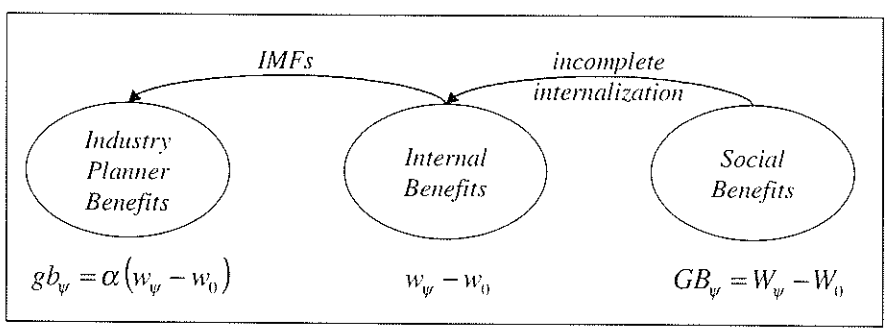

Figure 2: Schematic of Industry Planner Model

An important characteristic of the industry planner model is that prices are set at marginal costs and, at the same time, there are private incentives to innovate. In general, if firms can appropriate a fraction of the knowledge from their innovations, they will set prices above, and production below, competitive levels, so there should be deadweight loss[8]. The problem is most severe under constant returns to scale—an assumption I use throughout this paper—where marginal cost pricing means no profits and therefore no innovative rents!

In neglecting post-innovation monopoly behavior, I am neglecting a fundamental concern in theoretical innovation models from Arrow (1962) onward. The reasons is (1) that the advantages of simplifying and generalizing industry innovative behavior are worth this oversight and (2) the insights we derive from the simplified model would still be robust with the addition of a little deadweight loss.[9] For those readers requiring consistency with a particular microeconomic structure, the industry planner model represents a perfectly price-discriminating innovator that must limit price because of the potential for costless, partial imitation.

2.3 Net Welfare

Thus far, I have considered the socially optimal level of innovation and the level of industry innovation. This subsection discusses socially optimal *policy*, as opposed to optimal innovation, and defines the welfare measure, net welfare, for use in the models to come.

Let Z be a particular policy portfolio. Z might include the level of internalization, the type and level of technology-push measures, and the type and level of demand-pull measures. If the policy maker has some choice, perhaps being able to choose among technology policies but not among internalization policies, she will choose the policy that maximizes post-innovation welfare net of innovation costs[10]. The social planner's *policy* problem is written

[8]Innovator behavior depends on the sort of innovation and competitors' cost functions. If innovation decreases only inframarginal production costs, it might not alter market prices and quantities.

[9]It is an empirical matter whether deadweight loss is small relative to internal welfare gains. Deadweight loss is highest when (1) demand is elastic, (2) appropriability is high, and (3) there is an absence of price discrimination.

[10]Note that the net benefits of innovation and net welfare are *not* analogous concepts. The reason is that the net benefits of innovation are based upon a change in welfare whereas

mathematically as,

$$\max_{Z} \quad NW = W_{\psi(Z)} - \psi(Z) \tag{4}$$

where NW is **net welfare**. Note that $\psi(Z)$ is the *total* level of innovation; it includes both private and public efforts, given policy Z.

In the coming sections, I will compare NW under particular technology policies to first-best net welfare, NW_{fb}—net welfare under full internalization and socially optimal innovation.

2.4 Extensions to the Limited Model

The industry planner model as presented thus far is all that will be needed to understand the remainder of this paper. However, the industry planner is more flexible and versatile than has been presented thus far. In this section, I will discuss extensions and generalizations of the model that increase its scope of applicability. First I will show how intersectoral spillovers can be included in the model; next I will discuss how we might capture divergences in the social and private costs of innovation; and, last, I will discuss alternative approaches for capturing private underinvestment relative to internal benefits from innovation.

2.4.1 Spillovers

Interindustry spillovers are gains from innovation that accrue outside of the innovating industry, and that can be considered essentially unappropriable. For example, research into photovoltaics by British Petroleum may result in a better understanding of thin films that would be used by high-tech companies completely unrelated to photovoltaic manufacture. In contrast, the industry planner model presented thus far has considered only intraindustry spillovers—gains that accrue within the innovating industry and therefore might potentially be gained by the innovator. But interindustry spillovers are an important factor in the divergence between social and private rates of return from innovation. Hence, it would be nice to be able to include interindustry spillovers in the industry planner model.

---

net welfare is based only upon the final welfare level. Policy alters not only welfare with innovation, W_{ψ}, but also welfare without innovation, or W_0.

Including interindustry spillovers is a simple matter of tacking on an extra spillover benefit, which I will call $\gamma(\psi)$, to the social planner's problem, or

$$\max_{\psi} \quad NB_{\psi} = \overbrace{(W_{\psi} - W_0) + \gamma(\psi)}^{GB_{\psi}} - \psi \tag{5}$$

2.4.2 Divergence in Innovation Costs

The industry planner model presented thus far assumes that innovation costs for the two planner are identical: both are set at the real productive investment, ψ. In reality, though, both private and social innovation costs may differ from this ideal. Private costs may differ, for example, because of difficulties in raising funds for risky R&D ventures. Similarly, they may be lower than real costs because of R&D subsidies. Social costs may deviate from ψ through wasteful spending, redundancy in research effort, excess burden issues associated in raising public funds for innovation support, and so forth.

Let i_{ψ} and I_{ψ} represent private and social innovation costs associated with ψ. When these differ from ψ, the social planners problem and the industry planners problem can be written as

$$\max_{\psi} \quad NB_{\psi} = \overbrace{(W_{\psi} - W_0)}^{GB_{\psi}} - I(\psi) \tag{6}$$

$$\max_{\psi} \quad nb_{\psi} = \overbrace{\alpha(w_{\psi} - w_0)}^{gb_{\psi}} - i(\psi) \tag{7}$$

2.4.3 Alternative Models of IMFs

The Industry Planner Model thus far rolls all conventional IMFs into a multiplicative reduction in internal benefits, α. But the industry planner model is amenable to other approximations as well. Here I will discuss three of these.

Nordhaus (2002) implicitly provides the first approach. Instead of adjusting the private benefits of innovation, Nordhaus adjusts the private costs of innovation, ramping them up relative to the social costs, or

$$nb_{\psi}^{Z} = \overbrace{(w_{\psi}^{Z} - w_0^{Z})}^{gb_{\psi}^{Z}} - \eta\,\psi \tag{8}$$

where $\eta > 1$. This forces private actors to under-invest in innovation, just as in the approach above[11]

A second approach is to simply ramp up the discount rate associated with private sector innovation[12].

$$nb_{\psi}^{Z} = \overbrace{d\left(w_{\psi}^{Z} - w_{0}^{Z}\right)}^{gb_{\psi}^{Z}} - \psi \tag{9}$$

$$NB_{\psi}^{Z} = \overbrace{D\left(W_{\psi}^{Z} - W_{0}^{Z}\right)}^{GB_{\psi}^{Z}} - \psi \tag{10}$$

In other words, both planners discount the future, but the industry planner does so at a greater rate than the social planner, or $d < D$. Since the costs of innovation accrue before the benefits, this will both hold back private sector innovation and it will bias such innovation towards short-term projects.

A last approach springs from the new growth models of innovation and the approach of Goulder and Schneider (1997) and Goulder and Schneider (1999) in their intertemporal general equilibrium model of the U.S. economy. The idea is that the industry planner only perceives some of the gains from innovation as under his control, taking the rest as exogenous. This is captured functionally by,

$$W_{\psi}^{Z} = F(f(\psi), g(\psi)) \tag{11}$$

The social planner optimizes over ψ, but the industry planner optimizes only over $g(\psi)$, taking $f(\psi)$ as exogenous. If $\frac{\partial W}{\partial f(\psi)} \frac{\partial f(\psi)}{\partial \psi} > 0$, then the industry planner will invest below the socially optimal level.

I hope the reader can see that all of these approaches fall within the industry planner framework. They all assume a planner makes innovative decisions for the private sector, but that the planner's decisions are in some way deficient response to the internal gains from innovation.

---

[11] Nordhaus (2002) actually goes further than what I am presenting here, which is why I said that he "implicitly" demonstrates the approach. Nordhaus (2002) not only adjusts private costs of innovation, he also adjusts the social costs; his is a social planning model, so there is no distinction between the two. His intent is to capture not only under-investment given market prices for innovative services, but also that market prices lie below the true opportunity costs of innovative effort. In other words, the opportunity costs of research in one sector, the energy sector in Nordhaus (2002)'s model, may be higher then the market price if such research sucks research from other sectors where social returns are higher then private returns.

[12] I have not seen this approach before.

3 INNOVATION IN AN EET

This section section will explore policy, market innovative behavior, and socially optimal innovation in a two-period theoretical model of the development and emergence of an EET. Section 3.1 will introduce the model. Section 3.2 will discuss how the three policy options of interest—technology-push policies (broadly defined), adoption subsidies, and technology standards—are captured in the model. I will discuss the ramifications of incomplete internalization for policy and behavior in Section 3.3; the ramifications of IMFs in Section 3.4; and the combined effects of internalization and IMFs in Section 3.5. Section 3.6 will summarize the policy implications of the analysis. I will close, in Section 3.7, by discussing the implications of a particularly important model assumption.

3.1 The Model

3.1.1 Technology & Innovation Possibilities

The model to follow takes place over two periods. In the first period, actors—and this may include the government or private firms—invest in innovation[13]. In the second period, firms—and only firms—produce widgets into a homogeneous, competitive market. Demand for widgets is given by the constant elasticity function,

$$q = aP^{-b} \tag{12}$$

where a and b are parameters of the demand function, and P is the price of widgets. Since the market is competitive, we can also write $q = ac^{-b}$. The constant elasticity formulation gives downward sloping, convex demand, or $q' < 0$ and $q'' > 0$. The results in this section hold more generally for downward sloping, convex demand functions, but a specific functional form was necessary to obtain graphs in the results section.

Producing widgets requires two classes of inputs: standard inputs and environmental inputs. Standard inputs are those with costs internal to the

[13]Note that it does not matter which sorts of firms actually conduct the innovation. It could be the users of the technology, the producers of the technology, or a completely separate research industry. What matters is that all potential benefits from innovation make their way through the economic system to whomever might ultimately do the necessary research to improve widget production technology.

market absent government internalization policy; and environmental inputs are those associated with environmental consequences not captured in market prices absent internalization policy. For example, coal-fired electricity has standard inputs associated with fuel, maintenance, facility and transmission infrastructure, and so forth, and environmental inputs associated with carbon emissions, sulfur emission, and so on.

Widget production is constant returns to scale. Let x_π and x_ε represent the standard and environmental input requirements for one unit of output. These might vary depending on the production technology and the way that the production technology is operated. For a given pair, $\mathbf{x} = \{x_\pi, x_\varepsilon\}$, marginal private and social costs of production are given by,

$$c(x) = p_\pi x_\pi + \tau p_\varepsilon x_\varepsilon \tag{13}$$

$$C(x) = p_\pi x_\pi + p_\varepsilon x_\varepsilon \tag{14}$$

where p_π and p_ε are the opportunity costs, or prices, of standard and environmental inputs, and τ represents the level of internalization of otherwise external environmental costs[14]. If τ is zero, private costs include only standard costs. If τ is unity, internalization is complete, and private costs are equal to social costs.

Two fixed-proportions, constant-returns-to-scale production technologies can be used to produce widgets: TECHNOLOGY L, for "low-emissions", and TECHNOLOGY H, for "high-emissions". Figure 3 shows the two technologies in (x_π, x_ε) space prior to innovation. TECHNOLOGY L is the emerging technology: it has zero-emissions and higher standard costs than the dominant technology, TECHNOLOGY H[15]. TECHNOLOGY L might be electric cars[16] or PVs, and TECHNOLOGY H might be internal combustion engine vehicles or fossil fuel electricity generation.

As the arrow in Figure 3 shows, the standard input requirements of TECHNOLOGY L can be improved through innovation. Without innovation, the characteristics of TECHNOLOGY L are given by $\mathbf{x}_0^L = \{x_{\pi,0}^L, 0\}$.

[14] The mode of internalization is left unspecified. It might be emissions taxes, tradeable permits, or even the threat of stricter environmental standards

[15] The essential insights of the model still hold if TECHNOLOGY L's emissions are positive, but the intuition is most clear in the no-emissions case. I will discuss this issue in more detail in 3.7

[16] Note that electric cars may actually result in significant indirect emissions from the production of electricity.

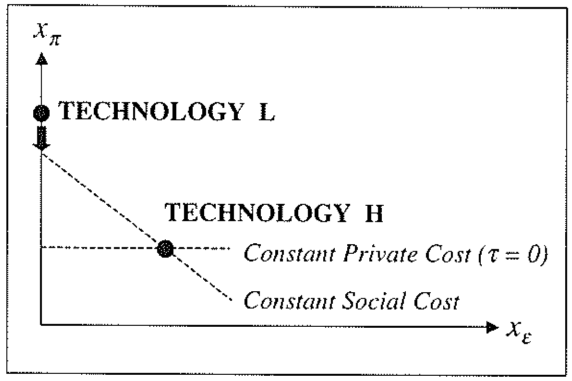

Figure 3: Innovation improves only TECHNOLOGY L.

With innovation level ψ, the characteristics of TECHNOLOGY L are given by $\mathbf{x}_\psi^L = \{x_{\pi,\psi}^L, 0\}$. (Recall that TECHNOLOGY L has zero emissions, so $x_\varepsilon^L = 0$.) The characteristics of TECHNOLOGY H are invariant at $\mathbf{x}^H = \{x_\pi^H, x_\varepsilon^H\}$.

Improvement in TECHNOLOGY L is given by the deterministic innovation production function,

$$x_{\pi,\psi}^L = x_{\pi,0}^L(1 - \mu_L\, \psi^{\nu_L}) \tag{15}$$

where $\mu_L > 0$ and $\nu_L \in (0, 1)$ are parameters. We might think of the innovation production function as representing a series of small discrete investment opportunities, each one less productive the previous; or we might interpret it as a set of unique innovation tracks, each leading a certain distance and having an associated cost.

The critical and fundamental characteristics of the innovation production function are (1) that innovation is beneficial—in other words, $x_{\pi,\psi}^L$ is decreasing in innovative effort, or $\frac{\partial x_{\pi,\psi}^L}{\partial \psi} < 0$—and (2) that there are decreasing returns to innovative effort, or $\frac{\partial^2 x_{\pi,\psi}^L}{\partial \psi^2} > 0$[17]. The model results and insights hold more generally for decreasing returns innovation production functions, but, again, a specific functional form was required to generate graphs in the results section.

[17]Note also that the marginal gains from innovation approach infinity as innovative effort approaches zero, or $\frac{\partial x_{\pi,\psi}^L}{\partial \psi} \to \infty$ as $\psi \to 0$. This implies an absence of fixed innovation costs.

The downward sloping line in Figure 3 is a social isocost line—social cost, C, is identical for all input combinations on this line. As Figure 3 shows, TECHNOLOGY L initially has not only higher private costs than TECHNOLOGY H, it also has higher social costs, or $C(\mathbf{x}_0^L) > C(\mathbf{x}^H)$. Hence, without innovation, it is best to produce with the dominant, more-polluting technology, even considering the full social costs of production. The horizontal line in Figure 3 is a no-internalization private isocost line—private cost is identical for all input combinations on this line when $\tau = 0$. This line will tilt upwards with increasing internalization, and, at full internalization, the social isocost line and the private isocost line will be identical.

The private isocost line through TECHNOLOGY H can be interpreted as an **adoption threshold**. It gives the input combinations at which TECHNOLOGY L can compete with TECHNOLOGY H on a private cost basis. Let the **threshold innovation level**, $\bar{\psi}$, be the level of innovation required for TECHNOLOGY L to reach the adoption threshold. Formally,

$$\bar{\psi} = \{\psi : c(\mathbf{x}_\psi^L) = c(\mathbf{x}^H)\} \tag{16}$$

It should be clear that the adoption threshold will become more severe the lower is internalization, or $\frac{\partial \bar{\psi}}{\partial \tau} < 0$. This point will be critical in interpreting the model.

Now let me introduce a final, and important, characteristic of the model. I assume that in the first-best world, a world with full internalization and no IMFs, it would be socially optimal would invest in TECHNOLOGY L to bring it past the adoption threshold and into the market. The private sector would undertake to do this because, in a first-best world, the private sector will act identically to the social planner. In other words, when $\tau = 1$ and $\alpha = 1$, there exists a $\psi > 0$ such that $NB_\psi = nb_\psi > 0$. Let the "first-best" innovation level be given by ψ_{fb}, and let net welfare at ψ_{fb} and under first-best conditions be given by NW_{fb}. What follows is an exploration of how to proceed when the world is not perfect—when there are IMFs and uninternalized environmental costs.

3.1.2 Welfare and the Net Benefits of Innovation

Now let us turn to the net benefits of innovation in the model. To do so, we must first look at welfare. For a given production technology, with

$\mathbf{x} = \{x_\pi, x_\varepsilon\}$, internal and social welfare under a constant-elasticity demand function are given by

$$w = \int_{q=0}^{q(c)} P(q)\,dq - c\,q(c) = \frac{a}{b-1}\,c^{1-b} + \Theta \tag{17}$$

$$W = \int_{q=0}^{q(c)} P(q)\,dq - C\,q(c) = \frac{a}{b-1}\,c^{1-b} - (1-\tau)\,p_\varepsilon x_\varepsilon\,aP^{-b} + \Theta \tag{18}$$

where Θ is some constant[18] [19].

Recall from Equation 2 and Equation 3, that NB_ψ and nb_ψ are based on change in welfare via innovation. Because TECHNOLOGY H is, by construction, the production technology absent any innovation, it is the basis for calculating welfare absent innovation, or $W_0 = W^H$ and $w_0 = w^H$. Welfare following innovation depends on which technology is used in the market following innovation, which depends on how far innovation reduces the standard input requirements of TECHNOLOGY L. If $\psi < \bar{\psi}$, firms will not adopt TECHNOLOGY L, so $W_\psi = W^H$ and $w_\psi = w^H$. In this case, innovation will be wasted: there will be no gross benefits, so $NB_\psi = nb_\psi = -\psi$. If, on the other hand, $\psi \geq \bar{\psi}$, firms will choose to use TECHNOLOGY L instead of TECHNOLOGY H, so $W_\psi = W_\psi^L$ and $w_\psi = w_\psi^L$. To summarize, the net benefits of innovation for the two planners can be written as,

$$NB_\psi = W_\psi - W_0 - \psi = \begin{cases} -\psi & \text{for } \psi < \bar{\psi} \\ W_\psi^L - W^H - \psi & \text{for } \psi \geq \bar{\psi} \end{cases} \tag{19}$$

$$nb_\psi = \alpha(w_\psi - w_0) - \psi = \begin{cases} -\psi & \text{for } \psi < \bar{\psi} \\ \alpha(w_\psi^L - w^H) - \psi & \text{for } \psi \geq \bar{\psi} \end{cases} \tag{20}$$

[18]Integrating the constant elasticity demand function becomes problematic when $b > 1$. The parameter Θ serves to eliminate the problem. The notion is that the constant elasticity formulation holds over all price ranges of interest, but at some high price demand drops to zero.

[19]Note that welfare here is intended to capture the full discounted benefits from the beginning of period two into the infinite future, but I have not explicitly accounted for this time dimension. For a more explicit accounting, we might multiply welfare by $\frac{1}{r}$, where r is an interest rate capturing time preference. Such a modification would have no meaningful effect on the results, so I have left it out for the sake of simplicity.

As per Equation 2 and Equation 3, the social planner will choose the innovation level that maximizes NB_{ψ}, and the industry planner, the proxy for market innovative behavior, will choose the innovation level that maximizes nb_{ψ}. Let $\tilde{I}$ represent the socially optimal innovation level. Let $\tilde{i}$ be the optimal innovation level for the industry planner.

3.2 Technology Policy Options

In the analysis to follow, I explore the performance of three technology policy instruments: technology-push instruments, technology standards, and adoption subsidies. Technology-push policies take place in the first-period of the model, during research and innovation. I assume that technology-push policies are akin to the social planner simply choosing the level of innovation. I further assume that technology-push policies are fully efficient—they bring about socially optimal innovation without any concomitant negative impacts, such as those associated with raising government funds or inefficiencies in the use of those funds.

Technology standards and adoption subsidies must act in the second period, after innovation, if they are to alter the private benefits of improving TECHNOLOGY L. Both of these demand-pull policies increase private benefits by changing the adoption threshold for TECHNOLOGY L. Technology standards eliminate the adoption threshold, so that $\bar{\psi} = 0$; adoption subsidies shift the adoption threshold, but do not eliminate it. Of importance, in the analysis to come adoption subsidies will never be able to achieve first-best net welfare because they reduce the cost of TECHNOLOGY L below its social cost—they introduce a price distortion. This does not mean they cannot increase net welfare[20].

3.3 Results: Incomplete Internalization

Now let us consider socially optimal innovation and private innovation when internalization is incomplete. Assume there are no IMFs, so we can isolate the effects of incomplete internalization. I will look at IMFs in 3.4 and

[20] An important difference between adoption subsidies and technology standards, but one that does not play out in this model, involves the distinction between prices and quantities under uncertainty (see Weitzman 1974). Adoption subsidies are a price instrument and technology standards are like a quantity instrument. Which is better under uncertainty will depend on abatement and environmental costs.

will combine IMFs with incomplete internalization in 3.5. I will leave the discussion of policy implications until 3.6.

The results in this section form a basis for everything to come. The essential idea I wish to get across here is that we can split "parameter space" into four possible regions, each with its own policy implications. These four regions will continue to hold when even when we add IMFs, and even when we allow for technological advance in both technologies in Section 4. For this reason, I will go into some detail in this subsection.

Figure 4 provides an introduction to the four regions. It shows socially and privately optimal innovation levels, $\tilde{I}$ and $\tilde{\imath}$, as a function of internalization *for one particular set of model parameters.* I have chosen the model parameters so that all four regions appear in Figure 4, but this need not be the case; under alternative parameters, only a subset of the regions might appear. Having all four, however, helps illustrate the relationships between the regions. Let us now proceed through the regions from right to left in Figure 4.

3.3.1 The "Hands Off" Region

The first region as we move to the left from full internalization in Figure 4 is the "Hands Off" region. This region represents a situation in which internalization has no effect on behavior and net welfare. Private firms invest to the first-best level in TECHNOLOGY L, this is socially optimal, and net welfare is first-best. Hence, the optimal policy is simply to let the market take its course. Two important questions are, first, why does the private

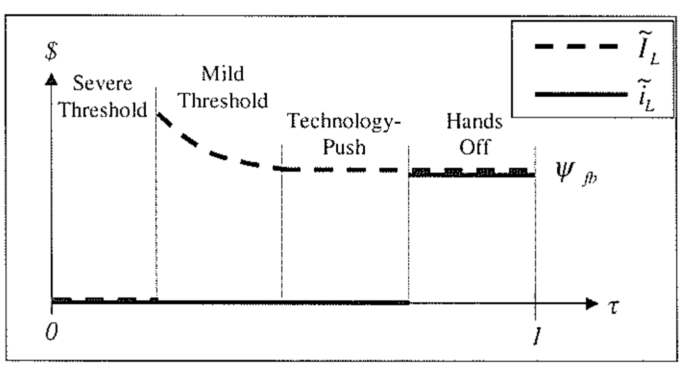

Figure 4: Optimal Innovation as a Function of Internalization

sector innovate to the socially optimal level and, second, why is net welfare first-best at this innovation level, even though internalization is incomplete?

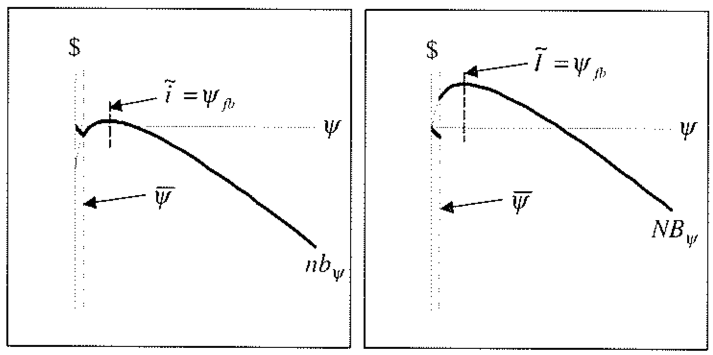

Industry Planner Net Benefits Social Net Benefits

Figure 5: The "Hands Off" Region

To answer these questions, and to understand the remaining regions, it is useful to investigate the net benefits of innovation in more detail. Figure 5 illustrates nb_ψ and NB_ψ in the "Hands Off" region. The thick, solid curve in the left panel is a plot of industry planner net benefits, nb_ψ, as a function of ψ. (Note that industry planner net benefits are simply internal net benefits because there are no IMFs.) The analogous thick, solid curve in the right panel plots social net benefits, NB_ψ, as a function of ψ. Note that NB_ψ is discontinuous at $\bar{\psi}$.

The long vertical dashed line in both panels represents $\bar{\psi}$. As discussed previously, net benefits for both planners are negative for $\psi < \bar{\psi}$. Hence, neither planner would choose a level of innovation less than $\bar{\psi}$ absent some sort of demand intervention to move the threshold—for example, a technology standard.

Right at $\bar{\psi}$, there are no private gains to innovation because $c(\mathbf{x}^L_{\bar{\psi}}) = c(\mathbf{x}^H)$, so $w_{\bar{\psi}} = w^H = w_0$ by construction. On the other hand, there are social gross benefits, $W_{\bar{\psi}} - W_0 = W_{\bar{\psi}} - W^H > 0$, because the switch from TECHNOLOGY H to TECHNOLOGY L reduces uninternalized environmental production costs. Whether $NB_{\bar{\psi}}$ is positive or negative, that is, whether the reduction in uninternalized costs outweighs the innovation costs at the adoption threshold, is a function of the model parameters. I have chosen parameters so that $NB_{\bar{\psi}} > 0$ in Figure 5. Regardless, NB_ψ will make a discontinuous jump upwards at $\bar{\psi}$ because of the discontinuous reduction in uninternalized environmental costs when firms stop using TECHNOLOGY H and begin using TECHNOLOGY L.

For $\psi > \bar{\psi}$, both planners reap positive gross benefits from innovation, or $w_\psi - w_0 > 0$ and $W_\psi - W_0 > 0$. In the "Hands Off" region, these gross benefits outweigh the innovation costs, so that $NB_\psi > 0$ and $nb_\psi > 0$, for a range of innovation levels. Hence, it is socially optimal to invest in TECHNOLOGY L, and the private sector will do just that. As ψ increases, nb_ψ and NB_ψ must eventually become negative because of decreasing returns to innovative effort.

Note the thin, solid curve to the left of the adoption threshold. This curve, combined with the thick, solid curve to the right of the adoption threshold, represents $W_\psi - W_0 - \psi$ for the social planner and $w_\psi - w_0 - \psi$ for the industry planner (recall that $\alpha = 0$) *under a technology standard*; that is, assuming that TECHNOLOGY L must be used in the market in the second period. This curve will prove important as we proceed. (Note that it does not truly represent *net benefits* under a technology standard, because these would use welfare with TECHNOLOGY L and no innovation, rather than welfare with TECHNOLOGY H, as a baseline. That is, net benefits under a technology standard would be given by $NB_\psi = W_{L,\psi} - W_{L,0} - \psi$ and $nb_\psi = w_{L,\psi} - w_{L,0} - \psi$, again assuming that $\alpha = 0$).

Now let us turn to why both planners invest to the first-best level and why this level of innovation achieves first-best net welfare. In words, if innovation is large enough to bring TECHNOLOGY L to market, then there will be no uninternalized environmental costs because TECHNOLOGY L produces no emissions. Hence, absent IMFs, there will be no market failures, and we will be in a first-best situation.

To understand the point more fully, and to understand the remaining regions, it is useful to go through the logic in more detail. Imagine that there will be a technology standard in the second period, mandating the use of TECHNOLOGY L. In this case, both planners will choose ψ so that the marginal benefits of innovation are equal to the marginal costs, or

$$\frac{\partial nb_\psi}{\partial \psi} = \alpha \frac{\partial (w_\psi^L - w_0^L)}{\partial \psi} - 1 = \alpha \frac{\partial w_\psi^L}{\partial \psi} - 1 = 0 \tag{21}$$

$$\frac{\partial NB_\psi}{\partial \psi} = \frac{\partial (W_\psi^L - W_0^L)}{\partial \psi} - 1 = \frac{\partial W_\psi^L}{\partial \psi} - 1 = 0 \tag{22}$$

Let $\hat{\imath}$ and $\hat{I}$ be the solutions to these marginal conditions.

Let me make three important points about $\hat{i}$ and $\hat{I}$. The first point is that $\hat{i}$ and $\hat{I}$ are independent of the baseline technology. Hence we get the same marginal conditions whether TECHNOLOGY L with no innovation is the baseline or whether TECHNOLOGY H is the baseline. This means that $\hat{i}$ and $\hat{I}$ are given by the innovation levels that correspond with the apexes of the combined thin and thick solid curves (thin to the left of $\bar{\psi}$ and thick to the right of $\bar{\psi}$) in Figure 5; these curves used TECHNOLOGY H as the baseline for calculating net benefits. The second important point is that $\hat{i} = \hat{I}$ absent IMFs. The reason for this is that TECHNOLOGY L has zero emissions. If the market must use TECHNOLOGY L in the second period—remember we are still working under the supposition that there is a technology standard—then there will be no environmental externality and $W = w$, which means the marginal conditions will be the same for the two planners. The third important point is that $\hat{i} = \hat{I} = \psi_{fb}$. The marginal conditions will be the same as in first-best; internalization is irrelevant if TECHNOLOGY L is to be used in the market in the second period.

Now we can answer our two questions. The fundamental characteristics of the "Hands Off" region are (1) that the marginal conditions fall beyond the adoption threshold, or $\hat{I} > \bar{\psi}$ and $\hat{i} > \bar{\psi}$, and (2) that net benefits are positive at the marginal levels, or $nb_{\hat{i}} > 0$ and $NB_{\hat{I}} > 0$. This means that socially optimal innovation and private innovation will be set at the marginal levels even without a technology standard, or $\tilde{i} = \hat{i}$ and $\tilde{I} = \hat{I}$. Recall from the discussion above that $\hat{I} = \hat{i} = \psi_{fb}$, so $\tilde{I} = \tilde{i} = \psi_{fb}$. Further, $NW_{\psi_{fb}}$ is independent of internalization because of the zero-emissions assumption, which implies that $NW_{\tilde{i}} = NW_{fb}$.

3.3.2 The "Technology-Push" Region

The next region as we move from right to left in Figure 4 is the "Technology-Push" region. This region captures a situation in which there is an EET that might get us out of our environmental problems, even under incomplete internalization, if the right level of research is devoted to improving it, but the private sector will not undertake the effort. This implies that technology-push policies can save the day.

Figure 6 illustrates net benefits in the "Technology-Push" region. Decreasing internalization has three effects on the system, altering it from the "Hands Off" region. First, the adoption threshold has moved to the right—it becomes more severe. Second, the private net benefits of innovation, nb_{ψ} have

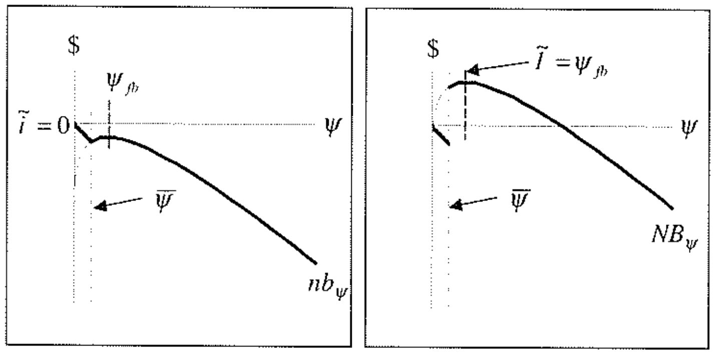

Figure 6: The "Technology-Push" Region

decreased for every ψ that still lies beyond the adoption threshold. Third, and conversely, NB_ψ has increased for every ψ that still lies beyond $\bar{\psi}$. (The effects of decreasing internalization on nb_ψ and NB_ψ are discussed in more detail in the appendix.)

What makes this region different functionally from the "Hands Off" region is that $nb_\psi < 0$ for all ψ. Hence, the private sector will stick with TECHNOLOGY H instead of pursuing TECHNOLOGY L, or $\tilde{i} = 0$, and this will not bring about first-best net welfare, or $NW_{\tilde{i}} = NW_0 < NW_{fb}$. At the same time, though, it is still socially beneficial to invest in TECHNOLOGY L to bring it past $\bar{\psi}$ because there exists a range of $\psi > \bar{\psi}$ for which $NB_\psi > 0$. As in the "Hands Off" region, and for the same reasons, $\tilde{I} = \hat{I} = \psi_{fb}$ and $NW_{\tilde{I}} = NW_{fb}$—getting innovation right will bring about first-best net welfare.

3.3.3 The "Mild Threshold" Region

The next region as we move from right to left in Figure 4 is the "Mild Threshold" region. This captures a situation in which the adoption threshold begins to constrain the socially optimal level of innovation. It is still net welfare enhancing to invest to bring TECHNOLOGY L to the point of adoption, but simply getting innovation right can no longer achieve first-best net welfare as it could in the "Technology-Push" region.

Figure 7 illustrates the situation in the "Mild Threshold" Region. The salient difference between the "Technology-Push" region and the "Mild Thresh-

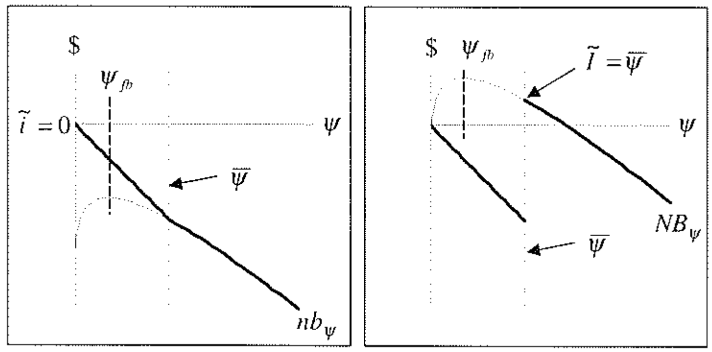

Industry Planner Net Benefits　　Social Net Benefits

Figure 7: The "Mild Threshold" Region

old" region is that the marginal innovation level lies to the left of the adoption threshold, or $\bar{\psi} > \hat{I} = \psi_{fb}$. In other words, to get the market to adopt TECHNOLOGY L, the social planner must *overinvest* in innovation—he must encourage more than the first-best level of innovation. It follows that $NW_{\tilde{I}} = NW_{\bar{\psi}} < NW_{fb}$.

The practical, real-world notion in the "Mild Threshold" region is that incomplete internalization can alter the adoption threshold enough to alter the socially optimal research program. It might, for example, call for a more risky program than otherwise, or it might call for a program that skips intermediate stages that would otherwise have benefits.

3.3.4 The "Severe Threshold" Region

The final next region as we move from right to left in Figure 4 is the "Severe Threshold" region. In this region, the adoption threshold is so severe that it is *socially detrimental* to bring TECHNOLOGY L to the point of adoption.

Figure 8 illustrates the situation. Social net benefits are negative for every level of innovation, so the it is socially optimal to abandon TECHNOLOGY L, which implies that $NW_{\tilde{I}} = NW_0 < NW_{fb}$.

The impact of incomplete internalization on socially optimal innovation behavior may seem surprising. We might expect that society should spend more on environmental innovation to make up for incomplete internalization: if we can't get prices right, then we should at least improve the technologies that we use in the market. The "Severe Threshold" region demonstrates the

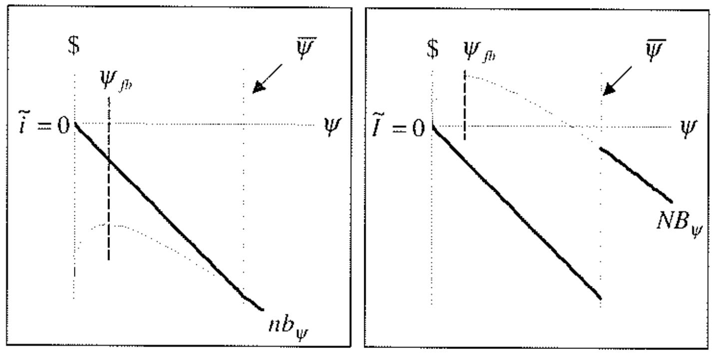

Industry Planner Net Benefits Social Net Benefits

Figure 8: the "Severe Threshold" region

possibility of the opposite effect: it may be that we should spend less on environmental technologies the lower is internalization because the adoption threshold has become too severe. For example, researchers may be confident that the ultimate technological potential of an EET lies beyond the full-internalization adoption threshold, but not the distorted adoption threshold under incomplete internalization. Proposition 1 formalizes this point.

Proposition 1 *Socially optimal environmental innovation may decrease with decreasing internalization when such innovation is embodied in an emerging environmental technology.*

3.4 Results: IMFs

Now let us reverse tack explore the effects of IMFs, assuming that internalization is complete, or $\tau = 1$. Figure 9, analogous to Figure 4, shows optimal innovation levels as a function of α. IMFs grow more severe as we move from right to left in the figure. As the figure shows, IMFs have two effects on $\tilde{i}$. (Note that socially optimal innovation is independent of IMFs.)

The first effect, as we move from right to left in the figure, is for $\tilde{i}$ to decrease with the severity of the IMF while still remaining positive—the private sector still chooses to invest in TECHNOLOGY L to bring it past the adoption threshold and into the market, but does so at a socially suboptimal level. This effect is due to a reduction in the marginal benefits of innovation for the industry planner from Equation 21.

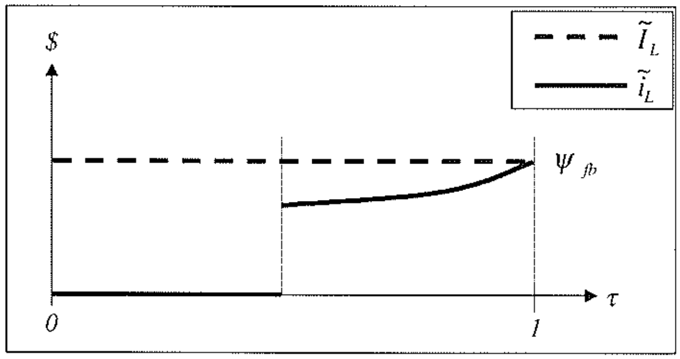

Figure 9: Optimal Innovation as a Function of IMFs

The second effect is a discontinuous change in behavior, just as was the case with incomplete internalization, with the private sector dropping innovation in TECHNOLOGY L altogether. This effect arises because IMFs change not only the marginal conditions for innovation, they also decrease the absolute gains from innovation. If IMFs are severe enough, gb_{ψ} will drop to the point that $gb_{\psi} < \psi$ for all ψ and the private sector will abandon TECHNOLOGY L.

3.5 Combined Effects of Incomplete Internalization and IMFs

Now let us combine the effect of IMFs and incomplete internalization. Figure 10 is a plot of optimal innovation as a function of internalization, analogous to Figure 4, only now we have assumed that $\alpha < 1$.

To start, note that the same four regions appear when we add IMFs to the mix. However, there are two important changes to regions. The first change is that the private sector will underinvest in innovation in the "Hands Off" region because IMFs reduce the private marginal benefits of innovation, as I discussed in the last section. The second effect is that the boundary between the "Hands Off" region and the "Technology-Push" region will shift to the right. It follows that the "Technology-Push" region is larger: it is more likely that the market will not pursue TECHNOLOGY L but doing so would achieve first-best net welfare. This second change is due to the reduction in the private gross benefits of innovation from IMFs, as I also discussed in the last section.

3.6 Policy Implications

We have now proceeded through a discussion of the effects of IMFs and incomplete internalization on private and socially optimal innovation behavior. It is now time to turn to policy. First, let me discuss optimal policy in each of the four regions, assuming that only one of the three policy approaches can be used. I will then discuss the benefits of combining technology-push policies with demand-pull policies.

3.6.1 Policy Implications by Region

In the "Hands Off" region, optimal policy is one of two choices: a hands-off policy or a technology-push policy. We know that the private firms will invest in the EET and bring it to market absent government intervention; but we also know that if IMFs are present—and they most certainly are in reality—the private sector will underinvest in the EET. This would indicate that technology-push policies are optimal because only they can induce the optimal level of innovation, thereby achieving first-best net welfare; a hands off policy cannot. To some degree, however, a goal of first-best net welfare in this region is philosophically troublesome. IMFs are broad-based: they will slow development not only of photovoltaic cells and fuel cells, but also toasters and high-definition TVs. If private firms will bring the EET to market, as they do in the "Hands Off" region, then one might argue that the goals of environmental technology policy have been satisfied and any further net welfare enhancements should be left to technology-policy more generally, and not environmental policy.

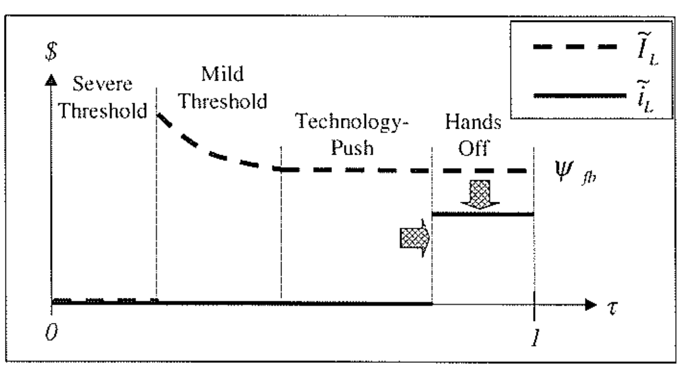

Figure 10: Optimal Innovation as a Function of Internalization with IMFs

The policy recommendation in the "Technology-Push" region follows its name. Technology-push policies will achieve first-best net welfare, and a hands off policy will result in no innovation whatsoever. This region lends support to policies such as the Partnership for a New Generation of Vehicles (PNGV) and the U.S. Department of Energy's Photovoltaic Program—programs aimed solely at increasing research in EETs. As with the "Hands Off" region, there is a legitimate question as to how far the government should push the technology in the presence of IMFs. One might argue that a reasonable policy would be to push the technology only to the point at which the private sector begins to pursue it of its own accord.

The outlook for technology standards and adoption subsidies in the "Technology-Push" region, and all other regions for that matter, is murky. While it is true that both policies can ensure the use of an EET in any region—technology standards by forcing the market to use the EET and adoption subsidies by providing financial incentives—their net welfare characteristics are unclear. Let me take the two policies in turn.

First, in the absence of IMFs, technology standards can, indeed, achieve first-best net welfare in all four regions; the private sector will invest to the private marginal innovation level in the presence of a technology standard and, from 3.3, we know that this is the first-best level and will result in first-best net welfare. In the presence of IMFs, however, technology standards will induce a socially suboptimal private sector innovation response. More problematic, technology standards cannot even ensure that the private sector will improve the technology to the point that it will be the low-social-cost option. To take a practical example, the California ZEV mandate will force ZEVs to be adopted, but it will not guarantee that auto makers will develop ZEVs to the point that it would be societally beneficial that they be adopted.

For the purposes of this paper, adoption subsidies can be thought of as a poor man's version of technology standards. They are similar to technology standards in that they can induce firms to use an EET, and are subject to the same concerns just expressed about the private response to technology standards. In addition, adoption subsidies introduce a price distortion by pushing the private cost of production using the EET below its true social cost. Hence, the best that adoption subsidies can do on a net welfare basis will always be lower than the best that technology standards can do. In short, then, adoption subsidies are subject to the same concerns as technology standards and more, in the Technology-Push" region and all other regions in which they might be used; technology standards will always be a superior choice[21].

Now let us move on to the “Mild Threshold” region. In this region, optimal policy will be one of two choices, neither of which can achieve first-best net welfare alone. The first choice is a technology-push approach. This approach cannot achieve first-best net welfare because it requires overinvestment in the EET. The second choice is technology standards or adoption subsidies. As before, however, we cannot be clear about the net welfare characteristics of these demand-pull subsidies technology standards because of IMFs, and, in the case of adoption subsidies, because of the welfare costs from the distortion in prices. Hence, the choice between technology-push and technology standards or adoption subsidies will ultimately be an empirical one in the “Mild Threshold” region.

In the "Severe Threshold" region, there are two choices. The first possibility is, again, technology standards or adoption subsidies; and, again, we cannot be sure of their net welfare characteristics for the same reasons as before. If these demand-pull policies are net welfare decreasing, then it will be optimal to pursue the second choice: a hands off policy.

3.6.2 Staged Technology Policy

Although individual technology policies may not always be able to achieve first-best net welfare, a combination of technology-push measures and demand-pull measures—technology standards in this model because of their superior net welfare characteristics—could always achieve first-best net welfare. Technology-push policies would be used to induce the optimal level of innovation, and demand-pull policies would be used to get the market to use EETs once they become the low-social-cost option. Because these two policy approaches take place sequentially, technology-push during the research phase and demand-pull following innovation, their combination implies a staged policy approach.

The advantage of the staged approach arises because, together, incomplete internalization and IMFs have two effects on private sector research; technology-push policies can target one and demand-pull policies can target the other. The first effect is a decrease private sector innovation in EETs. This arises both from IMFs and from incomplete internalization. The analysis thus far suggests that technology-push instruments should be

[21] In reality, however, adoption subsidies have advantages when demand is heterogeneous and technological advance is uncertain.

the primary means of attacking private underinvestment because demand-pull measures have two weaknesses in this regard[22]. The first weakness is that IMFs dampen the ability of demand-pull measures to increase private innovation incentives. The second weakness is implied simply through the model construction: demand-pull measures must take place *after* innovation to encourage innovation. This raises serious concerns about credibility and commitment[23]. For example, auto firms must believe that, even if they don't put any money into ZEVs, California regulators will force them to sell the product[24]. The second effect is a failure of private firms to adopt an EET once it becomes the low-social-cost option. This is clearly the realm of demand-pull policies.

3.7 The Zero-Emissions Assumption

Thus far, I have assumed that the EET creates zero environmental damages. This extreme assumption was meant to simplify intuition about, and to isolate certain important characteristics of, the development and emergence of EETs. In reality, however, most if not all technologies have associated with them some level of concomitant environmental damages. For example, photovoltaic cells often use nasty substances such as Gallium and Arsenic, and the production of almost any technology requires the use of electricity and other forms of energy which result in pollution or other environmental damages. How the results of the models in this section would change if TECHNOLOGY L had low, but not zero, emissions? Figure 11 shows the situation.

Figure 12, analogous to Figure 4 shows optimal innovation as a function

[22]There is a third compelling reason in the case of adoption subsidies. They are very difficult to target at potential customers in such a way as to efficiently increase the private potential for innovative rents.

[23]Advocates of demand-pull measures often point to learning effects from the use of technologies: putting them out in the field tells us something about their effectiveness and how they might be improved. In the context of the analysis here, if the programs are small, adoption subsidy programs like this are akin to R&D. That is, there is a distinction between limited programs to encourage limited use for information purposes, and large programs intended to create significant innovative rents, and therefore private incentive to innovate.

[24]I would like to acknowledge that the argument presented here is one-sided. I have discussed only the weaknesses of demand-pull policies without discussing the many weaknesses of government technology policies—for example, inefficiencies in raising government funds and supporting innovation that the private sector would undertake anyway.

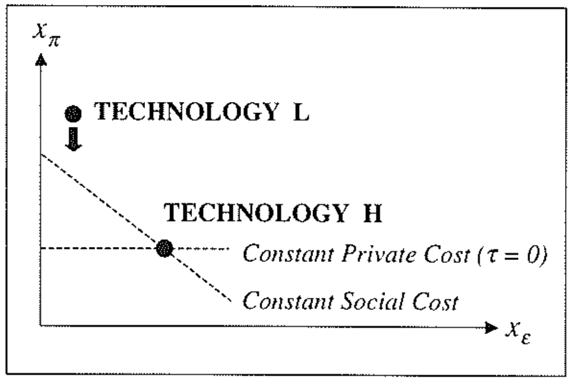

Figure 11: Positive Emission for TECHNOLOGY L.

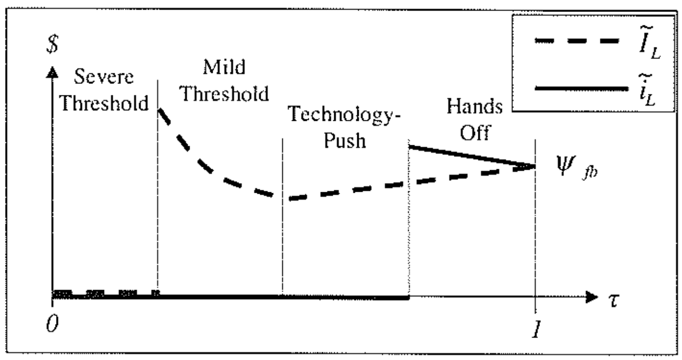

Figure 12: Optimal Innovation as a Function of Internalization

of internalization assuming no IMFs—IMFs will modify the figure results in the ways discussed previously: (1) shifting the boundary between the "Hands Off" region and the "Technology-Push" region and (2) decreasing private sector innovation in all regions that it is positive. I would like the concentrate here on three observations.

The first observation is that the same four regions are still in evidence when TECHNOLOGY L has positive emissions. Hence, they remain a useful intellectual organizing device. The second observation is that innovation investment in TECHNOLOGY L is a function of internalization for both planners in regions in which they choose to invest in TECHNOLOGY L. (I will leave the discussion of why innovation in polluting technologies might increase or decrease with internalization until Section 4).

The last observation is that the net welfare characteristics of technology policy are weakened in proportion to TECHNOLOGY L's emissions. Recall that some form of technology policy could always achieve first-best net welfare under the zero emissions assumption because once TECHNOLOGY L was used in the market there was no longer an environmental problem. If TECHNOLOGY L has positive environmental input requirements, we can no longer draw this conclusion. The degree to which the effectiveness of technology instruments are diminished by the presence of pollution in the EET is clearly an empirical matter, and one that I cannot resolve here. At the same time, it is fair to say that if the emissions characteristics of the emerging technology are small, as we might imagine them to be in the case of wind power or photovoltaics, we can probably feel comfortable with the net welfare and policy implications of the models in this section. On the other hand, if the emissions characteristics of the EET are significant, as one could argue they are for electric cars or fuel cell cars, then questions arise. A weaker, but more practicable, statement of the insights from this section might be that it is possible to achieve *close* to first-best net welfare using environmental technology policy. Even if we can fully achieve first-best net welfare, it is still useful to know what the best technology policy would be, and the results from this section still hold in this regard: optimal technology policy is a staged process beginning with technology-push measures and ending with demand-pull measures.

4 INNOVATION IN BOTH TECHNOLOGIES

Thus far, we have looked at a situation in which innovation could alter only the characteristics of the EET, TECHNOLOGY L. In this section, I will alter the simple model to allow for research in both the EET, TECHNOLOGY L, and the dominant technology, TECHNOLOGY H. My goal here is twofold. First, I wish to demonstrate the same four regions exist in this more complication model. Second, I wish to explore the effects of IMFs and incomplete internalization on the allocation of research investment between technologies.

4.1 The Model

The model for this section is shown in Figure 13. Firms might improve the standard costs of, say, PVs, or they might improve the characteristics of, say, combined-cycle natural gas plants or coal-fired power plants. Note that

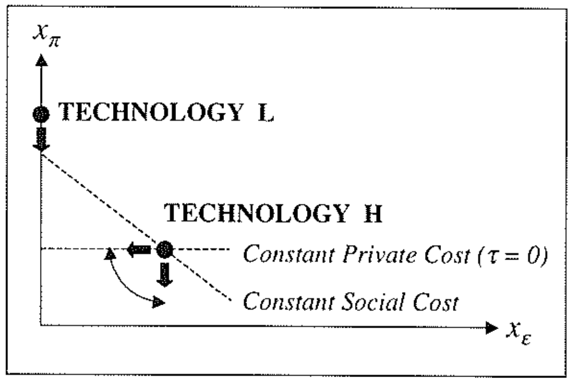

Figure 13: Innovation in Both Technologies

innovation in TECHNOLOGY H might fall anywhere from fully standard input reducing to fully environmental input reducing, as shown by the arc in Figure 13. For example, innovation in combined-cycle plants might make them cheaper to build, higher efficiency, or simply better in their emissions characteristics. Assume that the innovation production function for TECHNOLOGY H has the same functional form as that for TECHNOLOGY L but with different parameters (see Equation 15).

Assume, as well, that TECHNOLOGY L would the optimal choice in the absence of market failures, just as it was in Section 3. That is, in a perfect world, the market will bring TECHNOLOGY L to market rather than improve the performance of TECHNOLOGY H, and this would be the socially optimal strategy.

Let ψ_L and ψ_H be the research efforts in TECHNOLOGY L and TECHNOLOGY H. Mathematically, the decision situations for the two planners are now,

$$\max_{\psi_L,\psi_H} \quad nb_{\psi_L,\psi_H} = \overbrace{\alpha(w_{\psi_L,\psi_H} - w_0)}^{gb_{\psi_L,\psi_H}} - \psi_L - \psi_H \tag{23}$$

$$\max_{\psi_L,\psi_H} \quad NB_{\psi_L,\psi_H} = \overbrace{(W_{\psi_L,\psi_H} - W_0)}^{GB_{\psi_L,\psi_H}} - \psi_L - \psi_H \tag{24}$$

An important characteristic of the solutions to Equation 23 and Equation 24 is that the private sector will invest in one and only one technology, and the social planner will similarly invest in one and only one technology. The two planners invest in only one technology because, in this deterministic, homogenous setting, they know that only one technology will be used in the second period, and they know which one it will be through their research choices[25]. Innovation investments in the technology that will not be used in the market would be wasted. The reason they invest in at least one technology is as follows. Suppose that one of the planners has decided not to invest in TECHNOLOGY L. The planner must decide whether to invest in TECHNOLOGY H or to not invest in innovation at all. Because TECHNOLOGY H is the technology currently in use, any improvement will yield gross benefits. Since the production function for TECHNOLOGY H is such that the marginal improvements in technology approach infinity as investment approaches zero, it will always be optimal to invest at least some small amount to improve TECHNOLOGY H if an actor does not invest in TECHNOLOGY L[26]. Hence, it is always optimal to invest in either one or the other of the two research programs.

4.2 Results

4.2.1 The Four Regions

Figure 14 shows optimal innovation as a function of internalization when innovation reduces only the environmental input requirements of TECHNOLOGY H. Figure 15 shows optimal innovation as a function of internalization when innovation reduces only the standard input requirements of TECHNOLOGY H. These two scenarios span the range of innovation possibilities in Figure 13. In each figure, the top panel shows optimal innovation in TECHNOLOGY L, $\tilde{I}_L$ and $\tilde{i}_L$, and the bottom panel shows optimal innovation in TECHNOLOGY H, $\tilde{I}_H$ and $\tilde{i}_H$. Note that I have assumed no IMFs in the figures to facilitate intuition. IMFs will have the standard effects of (1) shifting the boundary between the "Hands Off" and the "Technology-Push" regions and (2) decreasing private innovation for all internalization levels.

As the figure shows, the same four regions are in evidence here, just as they were in Section 3. This is a fundamental point. It means that we can

[25]In reality, firms and governments invest many technologies at once to hedge against uncertainty and as a response to heterogeneity.

[26]The single exception to this is if $\alpha = 0$, in which case the market won't invest in any innovation.

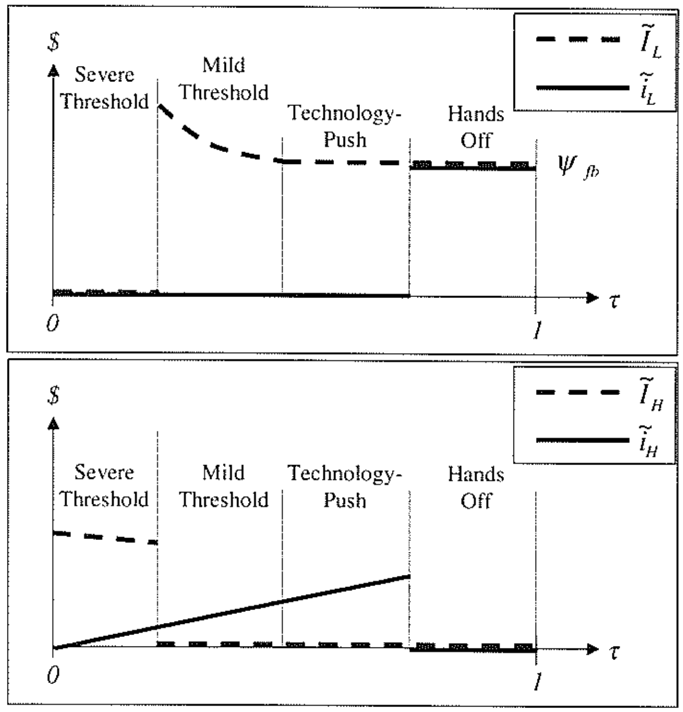

Figure 14: Optimal Innovation: Environmental Innovation in TECHNOLOGY H

still use the four regions as an intellectual organizing device.

There are, however, three important changes to the four regions. The first change is that the boundaries of the regions will be different than if we had assumed no opportunity to improve TECHNOLOGY H: for example, the "Hands Off" region will be smaller, and so forth. This is due to the fact that option to invest in TECHNOLOGY H will make TECHNOLOGY L relatively less attractive from both private and social perspectives; it must compete not with the pre-innovation state of TECHNOLOGY H, but rather an advanced version of TECHNOLOGY H. For example, PVs must compete not with the *current* state of fossil fuel technology, but with an advanced, future level.

The mathematical logic is as follows. Let $\hat{i}_H$ and $\hat{I}_H$ be the industry planner and socially optimal research investment in TECHNOLOGY H assuming a technology standard mandating its use. ($\hat{i}_H$ and $\hat{I}_H$ are analogous to $\hat{i}$ and $\hat{I}$ that we looked at in Section 3, only here we are referring to innovation in TECHNOLOGY H.) The market will invest to improve TECHNOLOGY L

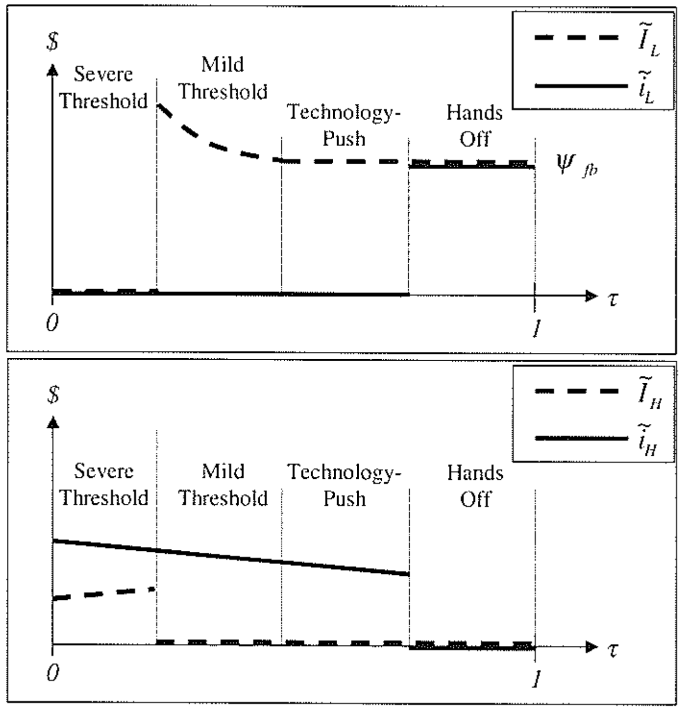

Figure 15: Optimal Innovation: Standard Innovation in TECHNOLOGY H

only if it can beat the best net benefits achievable from TECHNOLOGY H—if $nb_{\psi_L} > nb_{\hat{i}_H}$ for some ψ_L. Similarly, it is socially optimal to invest in TECHNOLOGY L if $NB_{\psi_L} > NB_{\hat{I}_H}$ for some ψ_L [27] [28].

The second important change from the previous section, as mentioned above is that the two planners will invest in TECHNOLOGY H when they do not invest in TECHNOLOGY L. Hence, the private sector will invest in TECHNOLOGY H in the "Technology-Push" region, the "Mild Threshold" region, and the "Severe Threshold" region, and the social planner will invest in TECHNOLOGY H in the "Severe Threshold" region.

[27] The criterion for the single program case in Section 3 was a special case of this multiple technology requirement in which $\hat{i}_H = \hat{I}_H = 0$, which implies that $nb_{\hat{i}_H} = NB_{\hat{I}_H} = 0$

[28] The adoption threshold is more complicated when innovation can improve TECHNOLOGY H. Innovating up to the point at which $c(\mathbf{x}_\psi^L) = c(\mathbf{x}_0^H)$ will no longer be good enough to get the market to adopt TECHNOLOGY L because there are private benefits to be made by bringing $c(\mathbf{x}_\psi^H)$ below $c(\mathbf{x}_0^H)$. If the government were to completely fund research in TECHNOLOGY L, the minimum level of innovation required get the market to adopt TECHNOLOGY L instead of investing in TECHNOLOGY H would be that for which $gb_{\psi_L} = nb_{\hat{i}_H}$.

The third important change is that the optimal level of innovation in TECHNOLOGY Hfor both planners is a function of internalization in those regions in which they invest in TECHNOLOGY H, whereas the same is not true for TECHNOLOGY L. When innovation primarily reduces the environmental input requirements of TECHNOLOGY H, as in Figure 14, then private innovation in TECHNOLOGY H will decrease with internalization, and socially optimal innovation will increase with internalization. Further, private innovation in TECHNOLOGY H will always be lower than socially optimal innovation in TECHNOLOGY H when both planners choose TECHNOLOGY H—in the "Severe Threshold" region. On the other hand, when innovation is primarily standard cost reducing, as in Figure 15, all of these results are reversed, although private innovation in TECHNOLOGY H may be lower than socially optimal innovation in TECHNOLOGY H in the "Severe Threshold" region if IMFs are severe enough[29].

The reason for these relationships is that TECHNOLOGY H has positive emissions, whereas TECHNOLOGY L does not. Because these results have been discussed in other papers[30], I will give only a cursory discussion here. In general, increasing internalization has two effects on private sector innovation in a polluting technology. The first effect, often referred to as the "direct" effect, is to increase the incentives for environmental innovation by increasing the level of private production cost reduction resulting from such innovation. The second effect, often called the "indirect" effect, is to decrease the incentives for any innovation because by decreasing a technologies market size. Hence, the private standard innovation will decrease with internalization, and private environmental innovation will increase or decrease with internalization depending on whether the direct or indirect effect dominates. In the model I have used here, the direct effect dominates.

The situation for the social planner is more complicated. To give the essential intuition, when internalization is incomplete, standard innovation

[29] Note that I have assumed that at the internalization level at which either planner switches from TECHNOLOGY L to TECHNOLOGY H, investment in TECHNOLOGY H is lower than that in TECHNOLOGY L. This seems the intuitive result, and is easy to obtain. It may, however, be possible to achieve the opposite result—that switching from TECHNOLOGY L to TECHNOLOGY H in response to decreasing internalization will *increase* aggregate innovation investment—under certain specification of the innovation production function. I will not explore this possibility in this paper.

[30] See Ulph (1997) for an excellent discussion also using a fixed-proportions technology. A short working paper on the subject is also available from the author.

has a negative side because it exacerbates over-production with a polluting technology. Hence, increasing internalization allows us to do more. Environmental innovation is even more beneficial under incomplete internalization than full internalization because it is more important to reduce the consequences of overproduction from the polluting technology. Hence, socially optimal innovation is decreasing in internalization[31].

Of great importance, none of these changes alter the fundamental policy implications from the last section. There is still a region in which it is best to simply let the market take its course; there are regions in which technology-push measures can improve net welfare; and a technology policy that stages technology-push and demand-pull policies will always be optimal.

4.2.2 The Allocation of Research Investment

Now we will turn to the allocative effects of IMFs and incomplete internalization. First I will present three propositions. I will then discuss their ramifications qualitatively. The essential point I wish to convey is that both IMFs and incomplete internalization can increase private innovation in dominant technologies, such as combined cycle natural gas facilities and internal combustion power trains. Further incomplete internalization can increase socially optimal innovation in these same dominant, more polluting technologies. The three propositions are as follows.

Proposition 2 *Socially optimal innovation in* TECHNOLOGY H *may increase in response to decreasing internalization, regardless of whether such innovation primarily reduces standard input requirements or environmental input requirements.*

Proposition 3 *Private innovation in* TECHNOLOGY H *may increase in response to decreasing internalization, regardless of whether such innovation primarily reduces standard input requirements or environmental input requirements.*

Proposition 4 *Private innovation in* TECHNOLOGY H *may increase in response to increasingly severe IMFs, regardless of whether such innovation primarily reduces standard input requirements or environmental input requirements.*

[31] See footnote 30.

Let me discuss these propositions in turn. The basis for Proposition 2 should be clear from Figure 14 and Figure 15. If decreasing internalization shifts the system from the "Mild Threshold" region to the "Severe Threshold" region, socially optimal innovation in TECHNOLOGY H will go from zero to something positive. In words, Proposition 2 says that it may be socially optimal to respond to decreasing internalization by shifting research funds away from, say photovoltaics, and towards improvements to fossil technologies; or from electric cars to improvements in traditional internal combustion power trains. Note that the result holds true even if innovation in TECHNOLOGY H is completely non-environmental, affecting only the standard input requirements of TECHNOLOGY H. This result is analogous to Proposition 1, and, as before, it runs counter to the notion that we should spend more on environmental innovation the lower is internalization.

Proposition 3 gives the analogous result for the private sector, and should also be clear from Figure 14 and Figure 15: decreasing internalization may shift the system from the "Hands Off" region to the "Technology-Push" region, thereby increasing private sector innovation in TECHNOLOGY H[32]. Note that this result is true even if innovation in TECHNOLOGY H reduces only the environmental input requirements of TECHNOLOGY H. In other words, automobile firms may be investing *too much* on reducing tail pipe emissions precisely because the environmental damages from these emissions have not been internalized.

Proposition 4 derives from the fact that IMFs shift the boundary between the "Hands Off" region and the "Technology-Push" region, just as they did in Section 3. Were the system to begin in the "Hands Off" region, with the market pursuing, say photovoltaics, then increasingly severe IMFs might force the market to abandon photovoltaics and increase innovation in fossil technologies.

Let me make three comments about Proposition 4. First, it is a result that holds outside of the environmental context of this paper, applying to any competition between a dominant and an emerging technology. IMFs bias innovation away from emerging technologies, not just emerging *environmental* technologies, and toward dominant technologies. We might therefore extend the result to a competition between two non-polluting technologies:

[32]As discussed above, the effect also occurs over the full range of private sector investment in TECHNOLOGY H—the "technology-push", "mild threshold", "severe threshold" regions—if innovation primarily reduces the standard input requirements of TECHNOLOGY H.

IMFs might shift the balance of private investment from renewables with small markets to those with larger markets, for example, away photovoltaics and into wind power. The second comment about Proposition 4 is that it runs counter to standard thinking about IMFs: that they decrease private sector innovation across the board. Proposition 4 tells us that IMFs have an important allocation effect, one that might feasibly result in *increased* private sector research, *but only in dominant technologies.* The last comment regarding Proposition 4 is that it did not require asymmetry in the effects of IMFs between technologies. We did not need to call on any "inertia" arguments—for example that firms are institutionally resistant to change and therefore biased against emerging technologies. All that is required is for IMFs to decrease the private gains from innovation.

The emerging character of the EET underlies all three of the above propositions. In all three cases, the market failures are causing innovation in the EET to drop to zero, thereby creating benefits from innovation in the dominant technology that would not otherwise be available. There are at least two interpretations of this effect. One is that there is a fixed cost to bring EETs to our homogenous market and the two market failures can drop the gross benefits below the fixed costs. IMFs reduce the gross benefits of innovation without altering the market threshold; incomplete internalization alters the market threshold and it changes gross benefits, reducing the gross benefits to private firms and increasing social gross benefits. In either case, the market failure renders research in the EET a net loser, so research in the EET is abandoned, leaving the market to the dominant technology, and inducing innovation in the dominant technology. This is a rather stark, all-or-nothing interpretation; the second interpretation is less extreme. Suppose that the two technologies compete in a heterogeneous market and the EET already has some demand. For example, photovoltaics are already in a number of niche markets. Improving photovoltaics will allow them to capture market share, but not the whole market, from their dominant rival, say fossil fuels. As the research in photovoltaics decreases in response to IMFs or incomplete internalization, it opens up profitable fossil technology research opportunities in the markets relinquished by photovoltaics. Hence, IMFs and incomplete internalization will not force abandonment of PVs; they will simply decrease innovation in PVs, with the same effect on fossil research as in the stark, homogenous interpretation.

The effect of IMFs on private sector innovation is particularly interesting in the heterogeneous case. On a first-order basis, IMFs will cause innovation to decrease in both PVs *and* fossil fuel technologies. IMFs will increase innovation in fossil fuels only if the second-order effect of PVs relinquishing market share dominates the first-order effect. The results in this section show that this will occur, even in the heterogeneous market, if there is a relatively abrupt adoption threshold.

One policy implication of this model with two possible research programs is that, in certain circumstances, *holding back* innovation in a dominant, polluting technology can open the way to socially beneficial innovation in EETs. This points to potential negative implications of government support of fossil fuel technologies, to take a practical example. Not only does such support consume a significant portion a limited national energy R&D budget, it reduces the benefits of innovation in EETs by increasing the perceived market threshold for these technologies.

5 CONCLUDING REMARKS

This paper has used a theoretical model to explore the emergence of EETs—for example, photovoltaic cells, fuel cells, and electric cars—and efforts to improve social welfare through government policies of support for these EETs. In this concluding section, I would like to highlight and discuss three important implications of this effort.

The first implication is that it is theoretically possible for environmental technology policy to achieve first-best, or close to first-best, net welfare. The essential requirements—requirements that were met in all the models I explored—are, first, that there exists a zero-emissions EET that can be improved, and, second, that in a first-best world, it would be optimal to improve this EET and bring it to market. Whether these conditions are met in a substantial number of real-world situations is an empirical matter.

The second implication is that environmental technology policy should be a staged process. Technology-push measures, such as research subsidies to private firms and research at government laboratories and universities, should be used to spur "pre-competitive" technological advance; demand-pull policies, such as adoption subsidies and technology standards, should be used if and when EETs are less costly to society then their dominant rivals. As I discussed in Section 3, a staged policy of this sort could always achieve first-best net welfare whereas policies that were limited to only one or the other could not in certain cases. The fundamental ideas behind this staging are that technology-push measures are the best approach to spur innovation and demand-pull measures are the only way to ensure adoption.

The last implication of the analysis is that the technological response to internalization policies may be highly nonlinear, both in the level of innovation and the allocation amongst technologies. This is an important issue when environmental damages are uncertain. Limited internalization policies, for example the recent proposal by Resources for the Future, may not serve as an effective hedge against bad climate outcomes: they may result not simply in too little environmental innovation, but essentially no innovation in the technologies that would be necessary to dramatically reduce emissions—technologies like photovoltaics and electric cars. These limited internalization policies will, instead, encourage the development of "middle scenario" technologies, such as higher efficiency cars and combined cycle natural gas facilities. This reinforces the need for technology policies to support the development of backstop technologies, serving as a hedge against very bad climate outcomes.

APPENDIX

The goal of this appendix is to demonstrate that, for $\psi > \bar{\psi}$, nb_ψ increases in internalization and that NB_ψ decreases in internalization in the model of Section 3. A secondary goal is to be more clear about the origin of these net benefits more generally.

Note that for any given ψ, the effect of internalization is only change gross benefits; it doesn't alter costs. Hence, this is where I will concentrate the discussion.

Figure 16 shows the demand function for TECHNOLOGY L in price-quantity space. Demand at $c(\mathbf{x}_0^L)$ is zero, because $c(\mathbf{x}^H) < c(\mathbf{x}_0^L)$ by assumption. If $\psi \geq \bar{\psi}$, then $c(\mathbf{x}_\psi^L) \leq c(\mathbf{x}^H)$ by construction, and the EET will take the whole market, which increases in size with further decreases in price through innovation.

First let us look at industry planner gross benefits. The internal gross benefits from innovation, $w_\psi - w_0$, are given by the shaded region between $c(\mathbf{x}^H)$ and $c(\mathbf{x}_\psi^L)$. It is only in movements beyond $c(\mathbf{x}^H)$ that the potential for private profits are created because the highest price any innovator could set for his product is the price of the next best technology, in this case

TECHNOLOGY H. Mathematically,

$$gb_{\psi} = \alpha(w_{\psi} - w_0) = \alpha \int_{c=c(\mathbf{x}_{\psi}^{L})}^{c(\mathbf{x}^{H})} q(c)\, dc \tag{25}$$

Increasing internalization raises the cost of TECHNOLOGY H. Formally, $\frac{\partial c(\mathbf{x}^{H})}{\partial \tau} = p_{\varepsilon} x_{\varepsilon} > 0$. It should be clear from Figure 16 that raising $c(\mathbf{x}^{H})$ will increase the internal gains from innovation, which will mean a rise in gb_{ψ}. Mathematically,

$$\frac{\partial gb_{\psi}}{\partial \tau} = \alpha q(c(\mathbf{x}^{H})) \frac{\partial c(\mathbf{x}^{H})}{\partial \tau} = \alpha p_{\epsilon} x_{\varepsilon}^{L} q(c(\mathbf{x}^{H})) > 0 \tag{26}$$

Now let us look at social gross benefits. GB_{ψ} differs from gb_{ψ} in that it includes not only the internal gains from innovation, but also the external environmental gains from switching from TECHNOLOGY L to TECHNOLOGY H. Mathematically, GB_{ψ} is given by

$$GB_{\psi} = \int_{c=c(\mathbf{x}_{\psi}^{L})}^{c(\mathbf{x}^{H})} q(c)\, dc + (1-\tau) p_{\epsilon}\, x_{\varepsilon}^{L}\, q(c(\mathbf{x}^{H})) \tag{27}$$

Increasing internalization changes GB_{ψ} by changing both internal benefits and the external environmental benefits. In total, the effect of increasing internalization is to decrease GB_{ψ}, or

$$\frac{\partial GB_{\psi}}{\partial \tau} = \overbrace{p_{\epsilon} x_{\varepsilon}^{L} q(c(\mathbf{x}^{H})) - p_{\epsilon} x_{\varepsilon}^{L} q(c(\mathbf{x}^{H}))}^{=0} + (1-\tau) p_{\epsilon} x_{\varepsilon}^{L} \frac{\partial c(\mathbf{x}^{H})}{\partial \tau} < 0 \tag{28}$$

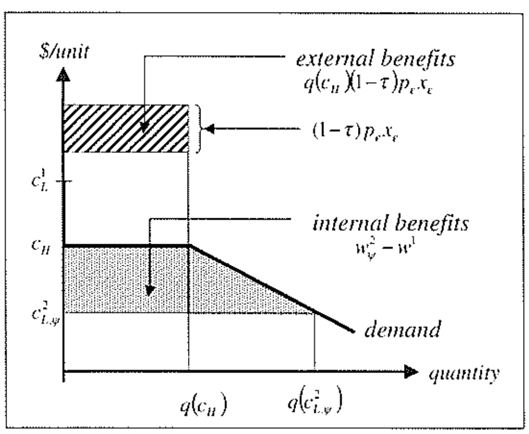

Figure 16: The Gross Benefits of Innovation in TECHNOLOGY L

In words, the change in internal benefits is exactly balanced by an opposite change that results from a change in the per-unit external benefits at constant q. In addition, though, increasing internalization decreases production with TECHNOLOGY H, thereby decreasing the benefits of switching from TECHNOLOGY H to TECHNOLOGY L at constant per unit external costs.

References

Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In R. Nelson (Ed.), *The Rate and Direction of Innovative Activity*. Princeton University Press.

Biglaiser, G. and J. Horowitz (1995). Pollution regulation and incentives for pollution control research. *Journal of Economics and Management Strategy 3*(4), 663–684.

Carraro, C. and A. Soubeyran (1996). Environmental policy and the choice of production technology. In C. Carraro, Y. Katsoulacos, and A. Xepapadeas (Eds.), *Environmental Policy and Market Structure*. Kluwer Academic Publishers.

Carraro, C. and G. Topa (1995). Taxation and enviromental innovation. In C. Carraro and J. Filar (Eds.), *Game Theoretic Models of the Environment*. Boston, MA: Birkauser.

Dosi, C. and M. Moretto (1997). Pollution accumulation and firm incentives to accelerate technological change under uncertain private benefits. *Environmental and Resource Economics 10*, 285–300.

Downing, P. and L. White (1986). Innovation and pollution control. *Journal of Environmental Economics and Management 13*, 18–29.

Duke, R. and D. Kammen (1999). The economics of energy market transformation programs. *The Energy Journal 20*(4), 15–65.

Fischer, C., I. Parry, and W. Pizer (1998). Instrument choice for environmental protection when technological innovation is endogenous. Discussion Paper 99–04, Resources for the Future, Washington, D.C.

Goulder, L. and S. Schneider (1997). Achieving low-cost emissions targets. *Nature 389*(4), 13–14.

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

Jung, C., K. Krutilla, and R. Boyd (1996). Incentives for advanced pollution abatement technology at the industry level: An evaluation of policy alternatives. *Journal of Environmental Economics and Management 30*, 895–111.

Katsoulacos, Y. and A. Xepapadeas (1996). Environmental innovation, spillovers and optimal policy rules. In C. Carraro, Y. Katsoulacos, and A. Xepapadeas (Eds.), *Environmental Policy and Market Structure.* Kluwer Academic Publishers.

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

Milliman, S. and R. Prince (1989). Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and Management 17*, 247–265.

Newell, R. (1997). *Environmental Policy and Technological Change: the Effects of Economic Incentives and Direct Regulation on Energy-Saving Regulation.* Ph.d. thesis, Harvard University, Cambridge, MA.

Newell, R., A. Jaffe, and R. Stavins (1998). The induced innovation hypothesis and energy-saving technological change. Discussion Paper 98-12, Resources for the Future, Washington, D.C.

Nordhaus, W. (2002). Modeling induced innovation in climate-change policy. In A. Grubler, N. Nakicenovic, and W. Nordhaus (Eds.), *Technological Change and the Environment.* Washington, D.C.: Resources for the Future.

PCAST (1999). *Powerful Partnerships: The Federal Role in International Cooperation on Energy Innovation.* The President's Committee of Advisors on Science and Technology: Panel on International Cooperation in Energy Research, Development, Demonstration, and Deployment.

Ulph, D. (1997). Environmental policy and technological innovation. In C. Carraro and D. Siniscalco (Eds.), *New Directions in the Economic Theory of the Environment.* Cambridge University Press.

Weitzman, M. (1974). Prices vs. quantities. *Review of Economic Studies XLI*, 477–491.