Profit-Maximizing R&D in Response to a Random Carbon Tax[*]

Erin Baker[†] and Ekundayo Shittu

University of Massachusetts, Amherst

November 19, 2004

[*]Running Head: Optimal R&D

[†]Correspondence Address:Erin Baker, 220 Elab, University of Massachusetts, Amherst, MA 01003; email: edbaker@ecs.umass.edu; 413-545-0670

Abstract

This paper determines a firm's profit-maximizing R&D response to an uncertain carbon tax, for two different R&D programs: cost reduction of low carbon energy technologies and emissions reductions of currently economic technologies. We find that optimal R&D does not increase monotonically in a carbon tax. R&D into alternative technologies increases only if the firm is flexible enough; R&D into conventional technologies first increases then decreases in a carbon tax. Firms that are very flexible may increase R&D into alternative technologies when the uncertainty surrounding a carbon tax is increased; otherwise firms will generally decrease R&D investment in uncertainty.

1 Introduction

A major topic in the discussion surrounding climate change policy is the role of induced technical change – technological innovation induced by climate policies. Another central issue when considering optimal policies for global climate change is the presence of uncertainty and learning about climate damages. There is much written about the impact of uncertainty on optimal abatement. The results from this literature indicate that the optimal response is to go slow,[1] and policy-makers in the U.S. seem to have taken this to heart, and are going slow indeed. Thus, there are no carbon-related policies to induce technical change at this time. Nevertheless, we see at least some firm-supported R&D: some firms are responding to the possibility of future policies.

In this paper we examine a firm's optimal R&D investment in response to a future uncertain carbon tax. A common assumption in the induced technical change literature ([17] [3] [10] [20]) is that a carbon tax will induce environmental R&D. Nevertheless, Sue Wing [23] has shown that this may not be true in a general equilibrium setting, and Farzin and Kort [5] show that the optimal level of environmental R&D can be non-monotonic in an emissions tax. We analyze these findings in a general micro-economic framework and find that they depend on the specific R&D program: all energy/environmental R&D is not alike. Moreover, a rational, forward-looking firm considering long-term R&D projects should base its decisions on beliefs about future policies, rather than simply react to current policies. Thus, understanding the role of uncertainty and learning is key to understanding firms' R&D investments. The motivation for this paper is threefold: 1) to provide guidance to firms on the best way to respond to uncertainty about climate policy; 2) to influence

[1] See, for example, [2][12][24][13]. The results are not, however, absolute. For example, both Webster [26] and Gollier et al. [8] show that the results can be reversed under certain condition.

the way that induced technical change is modeled in Climate-Economy models; and 3) to inform climate technology policy. The rationale for any technology policy is the gap that exists between the socially optimal investment in R&D and the privately provided R&D. This paper, through analyzing the impact of risk on privately provided R&D, is a first step in understanding the impact of risk on that gap.

We show that broad categories of technical change can be represented as changes in the relative prices the firm faces for its inputs. An application of Hotelling's Lemma then shows that the marginal benefits of R&D are proportional to well understood quantities. Namely, the marginal benefits of R&D into alternative technologies[2] are proportional to the unconditional demand for alternative energy inputs; the marginal benefits of R&D into reducing emissions of conventional, fossil-fuel technologies are proportional to the total carbon tax expenditures.

We find, for both programs, that the optimal investment into R&D can decrease with increases in a carbon tax. Investment into alternative technologies will increase monotonically with a carbon tax – if the firm is flexible enough. Investment into conventional technology will first increase and then decrease in a carbon tax, unless it is optimal to invest in carbon capture and sequestration of 100% of the emissions from a fossil fuel burning technology. Additionally, we analyze how an increase in the riskiness (in the Rothschild & Stiglitz sense of a mean-preserving spread [21]) of the carbon tax impacts optimal firm-level R&D spending. We find that, usually, R&D into conventional technology will decrease in uncertainty. On the other hand, R&D into alternative technology may increase in risk, if the firm is very flexible and if the increase in risk increases the probability that the carbon tax will be high enough to make non-carbon energy widely competitive. Otherwise,

[2] For example, programs to lower the cost of solar PVs or wind generation.

firms will decrease investment in alternative technologies as risk increases.

Most of the literature on induced environmental technical change uses an abstract abatement cost function, and makes assumptions about how technical change will impact this cost function ([4],[6],[10],[11],[15],[16],[18]). We illustrate, however, that even simple assumptions on how R&D impacts the production function have complex impacts on the marginal cost of abatement, impacts that depend on the characteristics of the firm[3] as well as on the description of the R&D program. Additionally, our framework allows for both input substitution and output reductions. Previous papers based on a production-function approach [5] have allowed for output reduction, but not substitution, while other papers based on a cost function approach [6] have allowed for substitution, but not output reduction. Evidence suggests that the majority of abatement results from input substitution [9]. Yet, at very high levels of abatement, it is likely that output reduction will play a role. When considering uncertainty, there is a positive probability of a very high tax (if damages turn out to be catastrophic) and thus it is important not to assume the possibility of reduced output away.

In the paper most closely related to this one Farzin and Kort [5] use a fully dynamic, optimal control model to consider the impact of both the timing and magnitude of an uncertain tax on the level of investment into abatement technology. They consider an investment in abatement technology that leads to a reduction in the emissions-to-output ratio and conclude that investment will increase up to a threshold tax, and then decrease; and that the presence of uncertainty discourages investment into abatement technology. The first result is driven by the fact that a very high tax will put downward pressure on output. The second result shows that, as long as the tax is smaller

[3] See Athey and Schmultzler [1] for a discussion of how firm flexibility impacts innovation.

than the threshold tax, optimal abatement under uncertainty is less than under certainty. On the other hand, Larson [14] provides an example, using a similar framework, where optimal abatement increases under uncertainty. Our model can explain these contrasting results by considering the impacts of general increases in risk. Additionally, our model generalizes the results in [5] by allowing for ex-ante substitution, corner point solutions, and by considering different types of R&D. Sue-Wing [23] uses a computable general equilibrium model to show that the negative output effect may cause aggregate investment into R&D to decline with an increase in the carbon tax.

In the next section we present the model and the logic of our analysis. In Section 3 we illustrate the impacts of R&D on the marginal cost of abatement. We then consider separately two different R&D programs: alternative technologies are discussed in Section 4 and conventional technologies are discussed in Section 5. Each of these sections is broken into 2 subsections. In Subsections 4.1 and 5.1 we analyze the optimal response of each R&D program to an expected increase in the carbon tax. In Subsections 4.2 and 5.2 we consider the impact of an increase in risk on each of the programs. In Section 6 we extend the analysis to energy efficiency programs. Section 7 concludes.

2 Firm-level decisions about R&D

We consider a hypothetical firm that both produces and uses R&D. The firm also produces carbon emissions as a joint product with output and thus is responsible for the carbon tax. An example of such a firm is American Energy Producers. Based out of Columbus, Ohio they are the largest electricity generator in the United States, and they have a significant R&D budget, including environmental R&D. We do not consider diffusion of R&D, thus this may be considered a model of R&D on firm-specific techniques.

We consider policies that impact the price of carbon such as a carbon tax or a cap-and-trade policy. We assume that there is no price on carbon currently, but that – depending on both scientific and policy variables – there may be a carbon tax in the future. For this study we assume that the probability distribution of the carbon tax is independent of the investment in R&D – the firm is a "policy taker." We assume that decisions about R&D must be made now, so that the changes in the production function can be implemented when the actual carbon tax is realized, and that the returns to R&D are deterministic.

We focus on two broad environmental R&D programs. The first is a program aimed at reducing the cost of alternative, non-carbon energy technologies, such as wind and solar. The second program is aimed at reducing the emissions of conventional, fossil-fuel-burning energy technologies. We use a profit function approach within a two-period theoretical model. Investments in R&D are made in the first period under uncertainty about a carbon tax. For simplicity, we ignore production in the first period. Optimal production is chosen in the second period after the firm learns about the carbon tax and technical change has been achieved, leading to a second period profit function. We take the market structure for output to be exogenous – the firm either faces a competitive price for its output or faces a known downward sloping demand curve.

Consider a firm that uses three inputs – non-energy inputs x, carbon energy inputs, e_c and non-carbon energy inputs, e_{nc}. Let e_c be normalized so that, using the current technology, one unit of carbon energy produces one unit of emissions. Then the total firm-specific price paid for carbon energy is the cost of the fuel, p_c, plus the price of the carbon emitted, t. Assume that, under the current technology, non-carbon energy is more expensive than carbon energy: the firm-specific price of non-carbon energy equals $p_c + \eta$. The price of non-energy inputs is w. Finally, let α represent the amount of technical change. Thus, the firm's second period profit function, assuming that the carbon tax, t, is known, is $\pi\left(w, p_c + t, p_c + \eta; \alpha\right)$.

The impact of a carbon tax on the firm's behavior depends on the firm's flexibility to substitute between carbon and non-carbon energy as measured by the elasticity of substitution. Surprisingly, there is little evidence on this topic. Popp [19], in his climate economy model including endogenous technological change, has calibrated the short-term elasticity of substitution between carbon and non-carbon energy as 1.6, implying that they are substitutes, but not very close substitutes. On the other hand, if we consider electricity generators, then in the long run, carbon and non-carbon energy are perfect substitutes. The rigidities inherent in the long lifetimes of generation plants determine the medium-term elasticity. Thus, we will present our results under different assumptions about the elasticity of substitution.

Let $g\left(\alpha\right)$ represent the cost of technical change, and assume that it is increasing and convex. The firm chooses the level of α in the first period, when the carbon tax is still unknown. In the second period the firm learns the value of the carbon tax t and maximizes profit accordingly. The firm's two period problem is:

$$\max_{\alpha} -g\left(\alpha\right) + \mathrm{E}_t\, \pi\left(w, p_c + t, p_c + \eta; \alpha\right) \tag{1}$$

The first order condition for α is

$$g'\left(\alpha\right) = \mathrm{E}_t\left[\frac{\partial \pi}{\partial \alpha}\right] \tag{2}$$

The question we ask is, how does the optimal amount of R&D spending change with changes in the probability distribution of t? If the right-hand side of (2) increases, then the optimal level of R&D spending increases. Thus we will focus on how a change in the probability distribution of t impacts the expected marginal increase in profits due to an increase in R&D. In Subsections 4.1 and 5.1 we consider increases in the expected value of t, while in Subsections 4.2 and 5.2 we consider a mean-preserving spread of t. We consider two categories of technical change – R&D aimed at reducing the cost of no-carbon alternatives, and R&D aimed at reducing emissions of conventional fossil-fuel technologies.

2.1 R&D into alternative technologies

First, we consider R&D into reducing the cost of non-carbon alternatives (denoted ARD for *Alternative technology R&D*). We model ARD as reducing the premium η for non-carbon energy by a fraction α. Thus the second period profit function is

$$\pi\left(w, p_c + t, p_c + (1 - \alpha)\,\eta\right) \tag{3}$$

This program could represent, for example, a firm's research into improving the efficiency of their wind turbines.

2.2 R&D into conventional technologies

In this section we consider R&D into reducing emissions of the currently economic technology (denoted CRD for *Conventional technology R&D*). We model CRD as reducing the carbon intensity of a unit of carbon energy from 1 to $(1-\alpha)$. Thus, effectively, the price of carbon energy is reduced from $p_c + t$ to $p_c + (1-\alpha)\,t$. This program represents an investment into technology that will capture a fraction α of the firm's carbon emissions. In Section 6 we show that an R&D program into increasing the efficiency of carbon energy has a qualitatively similar response to a carbon tax as CRD. Thus, the results of this section can be applied more broadly.

Under CRD the second period profit function is

$$\pi\left(w, p_c + (1-\alpha)\,t, p_c + \eta\right) \tag{4}$$

Below, we show how these R&D programs impact the marginal cost of abatement. We note that different R&D programs impact the marginal cost of abatement in different ways.

3 Marginal Abatement Cost curves.

In this section we discuss how the above programs, aimed at particular technologies, impact the Marginal Abatement Cost curve (MAC), and show that the impact of technological change on the MAC is complex: it depends on the characteristics of the production function and on the R&D program.

Figure 1 illustrates the impact of alternative R&D on the MAC under three assumptions about the elasticity of substitution between carbon and non-carbon energy. The x-axis shows abatement, defined as the fraction of emissions reduced below a business-as-usual optimal. The y-axis is the marginal cost of abatement. The MAC is important because firms will choose their level of abatement so that the marginal cost of abatement is just equal to the carbon tax. The solid line represents the MAC before R&D, and the dotted line represents the MAC after R&D into alternative technology. In the first panel, carbon and non-carbon energy are perfect substitutes. Abatement is achieved entirely through output reduction, until the stringency of abatement (about 70% in this example) makes it profitable for the firm to begin to switch to non-carbon energy. At this point the MAC is flat, as the firm substitutes one unit of non-carbon energy for one unit of carbon energy to achieve each additional unit of abatement. The dotted line shows that R&D has no impact on lower levels of abatement, but impacts the point at which the firm switches over to non-carbon energy. The next panel shows a firm with high, but imperfect, substitutability. In this case, R&D has an increasing impact on the MAC, up until about 90% abatement. At that point, due to imperfect substitutability, the marginal cost of abatement becomes dependent on the cost of reducing the last bit of carbon energy. Finally, the third panel shows the special case where carbon and non-carbon are complements. In this case, since there is no substitution, abatement is achieved entirely through output reduction. Reducing the cost of either of the inputs increases the marginal cost of reducing output: the MAC is increased by R&D. Note that this does not imply that performing R&D is irrational in this case. What matters to the firm is not the marginal cost of abatement, but the total profits. Reducing the cost of an input is always (weakly) beneficial to the firm.

Figure 1 clearly shows that the impact on the MAC is complex: in none of these cases do we see a proportional reduction that retains the shape of the original MAC.

Figure 2 illustrates the impact of R&D into conventional technologies on both the cost of abatement and the MAC. The left hand panel illustrates the impact of a 50% reduction in emissions on the cost of abatement. Note that the cost of abatement is everywhere (weakly) lower after R&D;

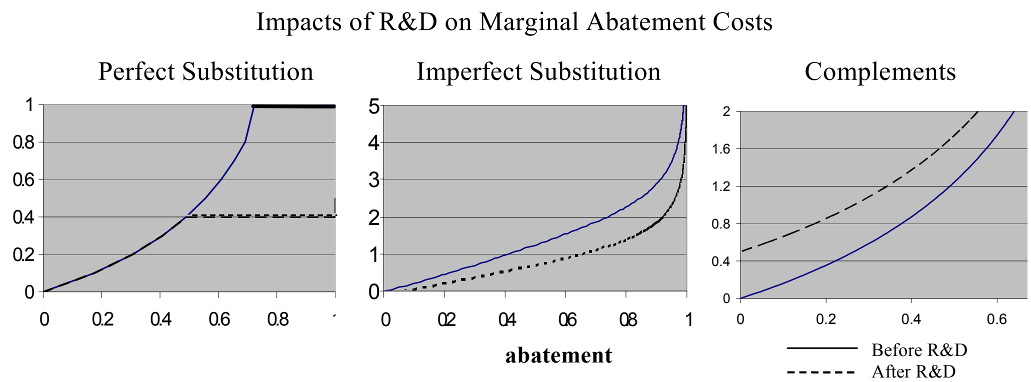

Figure 1: The impact of Alternative R&D on the marginal cost of abatement. The first panel represents a firm with perfect substitution between carbon and non-carbon energy; the second panel a firm with high, but imperfect substitution; the third panel, a firm for whom carbon and non-carbon energy are complements.

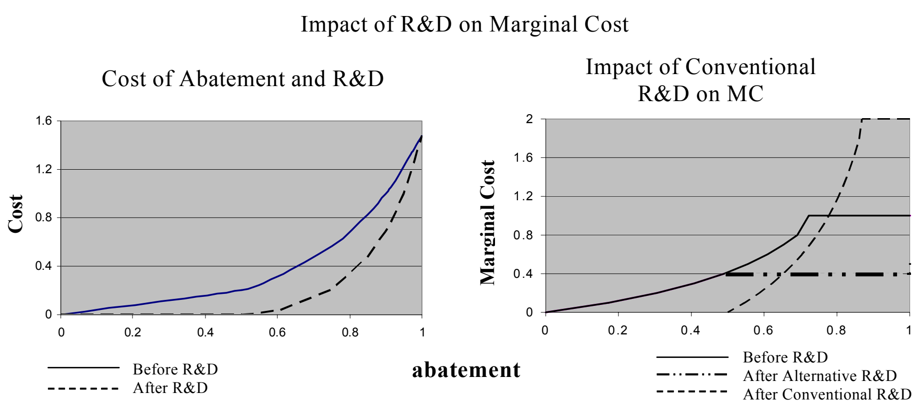

Figure 2: The left-hand panel shows the cost of abatement before and after conventional R&D. The right-hand panel shows the marginal cost of abatement in three cases: before R&D, after conventional R&D, after alternative R&D.

however, the cost of full abatement remains unchanged. This is because in order to achieve full abatement the firm must stop using the conventional, fossil-fuel burning technology. Thus, the slope of the cost curve is steeper after R&D over the higher range of abatement. Since the firm has improved their conventional technology, the jump from conventional to alternative energy is now greater. Thus, as shown in the right hand panel, the *marginal* cost of abatement is greater after R&D for abatement above 70%. This phenomenon will exist whenever firms take a ladder approach to abatement, substituting progressively cleaner technologies for dirtier ones as restrictions get tighter.[4] An investment into one technology on the ladder will increase the marginal cost of abatement for higher levels of abatement, since higher levels of abatement will require the firm to substitute away from the improved technology. The qualitative impacts of conventional R&D are similar under different assumptions about substitutability.

For illustration, the right-hand panel of Figure 2 also shows the impact of alternative R&D on the same baseline MAC, highlighting the difference between the two programs. CRD has a big impact on lower levels of abatement, but increases the marginal cost of high levels; ARD has a small impact on low levels of abatement, but a large impact on higher levels.

4 R&D into Alternative Technologies

Under ARD the right-hand side of the first order condition (2) simplifies to:

$$\mathrm{E}_t\left[\frac{\partial \pi}{\partial \alpha}\right] = -\eta\,\mathrm{E}_t\left[\frac{\partial \pi}{\partial p_{nc}}\right] \tag{5}$$

[4] See [7] for an example of sulpher dioxide abatement.

where p_{nc} represents the total price of non-carbon energy. Applying Hotelling's lemma, the first order condition (2) simplifies to

$$g'(\alpha) = \eta \, \mathrm{E}_t \left[e_{nc}^* \right] \tag{6}$$

where e_{nc}^* represents the unconditional demand for non-carbon energy. The marginal benefit from ARD is perfectly proportional to the demand for non-carbon energy. A firm will invest in alternative technology to the degree that they expect to be using alternative technology in the future. This simple and intuitive result allows us to analyze the unconditional demand for non-carbon energy and map the results back onto the optimal investment in R&D. If demand for non-carbon energy increases in the carbon tax, this implies that optimal investment in ARD increases in the expected value of the carbon tax.[5] Similarly, if the expected unconditional demand for non-carbon energy increases in uncertainty, then optimal investment in ARD will increase in uncertainty.

4.1 Increase in the expected value of a carbon tax: ARD

How will the demand for non-carbon energy respond to an increase in the price of carbon energy? This depends on two, opposite, effects. On the one hand, *conditional* demand (holding output constant) for non-carbon energy increases in the price of carbon energy; this is the substitution effect. On the other hand, output decreases in the price of an input. Which of these effects dominates depends on the relative values of the elasticity of substitution between carbon and non-carbon energy and the price elasticity of demand the firm is facing. The higher the elasticity of substitution, the more likely that total demand for non-carbon energy will rise in the carbon tax.

[5] Assuming an interior solution. We show below, however, that if the interior solution is monotone increasing, then the overall optimal solution is also increasing.

On the other hand, a higher price elasticity of demand for the firm's output makes it more likely that non-carbon energy will decrease in the carbon tax.

In fact, we show below, that for a monopolist with a nested CES production function facing a CES demand curve, R&D will unambiguously rise in the carbon tax if the elasticity of substitution between carbon and non-carbon energy is higher than both the elasticity of demand and the elasticity of substitution between energy and non-energy.

4.1.1 Constant elasticity production and demand

In this section we consider a monopolist, but the results are exactly the same for a competitive equilibrium with the same production function (see author for details). Consider a monopoly with a constant elasticity demand ($Ay^{-\frac{1}{b}}$) and a nested CES production function:

$$f(x, e_c, e_{nc}) = \left[x^{\rho} + (e_c^{\gamma} + e_{nc}^{\gamma})^{\frac{\rho}{\gamma}}\right]^{\frac{1}{\rho}} \tag{7}$$

See the Appendix for the proof of all propositions.

Proposition 1 *If the elasticity of substitution between carbon and non-carbon energy is higher (lower) than both the elasticity of substitution between energy and non-energy and the price elasticity of demand, then R&D into reducing the cost of non-carbon energy increases (decreases) in the carbon tax.*

The case where the elasticity of substitution between carbon and non-carbon energy is lower than both the elasticity of substitution between energy and non-energy and the price elasticity of demand corresponds with the case where the MAC is increased by R&D, and thus it is not surprising that the optimal investment into R&D decreases in a carbon tax. This case, however, is most likely not widespread. For most firms, investment into ARD will increase in an expected carbon tax. The impact of output reduction is tempered by the ability to substitute. Note that this result differs from that in Farzin and Kort: there is no threshold tax above which ARD decreases.

4.2 Mean-preserving spread in carbon tax: ARD

In this section we determine the impact of an increase in risk on the optimal level of ARD, focussing on firms that increase investment in R&D with an increase in the carbon tax, i.e. firms who have a high enough elasticity of substitution between carbon and non-carbon. Recall that any change in the distribution of t that increases the right-hand side of (2) will increase the optimal level of R&D. The expected value of a function increases in risk if the function is convex and decreases if the function is concave.[6] If the function is neither convex nor concave, then it will increase with some increases in risk and decrease with other increases. We use this framework to analyze how the demand for R&D will respond to an increase in the risk of the carbon tax.

In order to determine how the optimal level of R&D changes with an increase in risk, we need to determine if e_{nc}^* is convex or concave in t. If e_{nc}^* is concave in t then expected demand, $\mathrm{E}_t\left[e_{nc}^*\right]$, decreases in risk, thus α^*, the optimal amount of R&D, decreases in risk. If e_{nc}^* is neither concave nor convex it means that the optimal level of R&D decreases with some increases in risk, and increases with other increases in risk.

While theory gives us some information about how unconditional demand responds to an in-

[6] Jensen's inequality is a well-known special case of this fact. See Rothschild and Stiglitz [21] and [22] for more details.

crease in a cross price, it has little to say about the convexity or concavity. We can, however, say this: unconditional demand for an input, as a function of its cross price, is bounded above. As t increases, output will decrease, or in the extreme case, stay the same. Since we are assuming that the inputs are substitutes ($\sigma > 1$), we can say that for any finite output, the required input must be finite. In particular, for any level of output $\tilde{y}$ there exists $\tilde{e}_{nc}$ such that $f(\tilde{x}, 0, \tilde{e}_{nc}) = \tilde{y}$. Thus, since e^*_{nc} is increasing but bounded from above, it must be concave for large t. Proposition 2 makes this precise (See appendix for proof).

Proposition 2 *Assume that optimal R&D into alternative technology increases in the carbon tax. Then, optimal R&D into alternative technology will decrease with some increases in risk.*

More specifically, it can be shown (see author for details) that there exists a tax $\hat{t}$ such that any Mean-Preserving Spread (MPS) skewed to the right of $\hat{t}$ will decrease the optimal investment in alternative R&D. This has two implications. First, firms that believe that the main threat of climate change is a very low probability of a very high carbon tax should not invest greatly in technology R&D. Second, a misplaced focus on low-probability, high value events by policy makers or environmentalists may be counter-productive in terms of technology R&D investments. A balanced view of climate change and the possible tax implications will lead to a higher level of R&D.

Is e^*_{nc} always concave in t? No, it depends again on the elasticity of substitution. Consider a production function that is linear in carbon and non-carbon inputs. As long as $t \leq \eta$ the firm will use no non-carbon inputs; when $t > \eta$ the firm will jump to using all non-carbon inputs, and the demand will be constant in t thereafter. Thus, for some elasticities of substitution, e^*_{nc} is neither convex nor concave, implying that optimal R&D is ambiguous in risk. To illustrate this further we present a continuation of the example from Section 4.1.1.

4.2.1 Constant elasticity production and demand

The left panel of Figure 3 illustrates the demand for non-carbon energy as a function of t for three different elasticities of substitution. We assume that the price elasticity of demand is close to 1 and the elasticity of substitution between energy and non-energy is .75. The qualitative results illustrated in the figure are not sensitive to either of these assumptions. The elasticity of substitution between carbon and non-carbon energy, however, has a significant qualitative impact, as can be seen from the figure.

The figure shows that when carbon and non-carbon energy are not close substitutes, then demand for e_{nc}^* is everywhere concave, and thus R&D decreases unambiguously in risk. When, however, the firm is more flexible, the demand for e_{nc}^* is convex and then concave; thus increases in risk may increase or decrease innovation. Note that the inflection point, $t = 1$ on this figure, is the point where non-carbon energy becomes economically competitive with carbon energy, i.e. where $p_c + t = p_c + (1 - \alpha)\,\eta$.

For intuition, consider a firm that has perfect substitution between carbon and non-carbon energy. This firm will use only carbon energy until the carbon tax increases to the point that the cost of carbon and non-carbon energy are equal. At that point, the firm will switch to all non-carbon energy, and the carbon tax will have no additional impact. Thus, if the firm expected a small carbon tax, they would do no R&D; if they expected a large carbon tax they would do the maximal amount of R&D. A change in the probability distribution of the carbon tax will induce

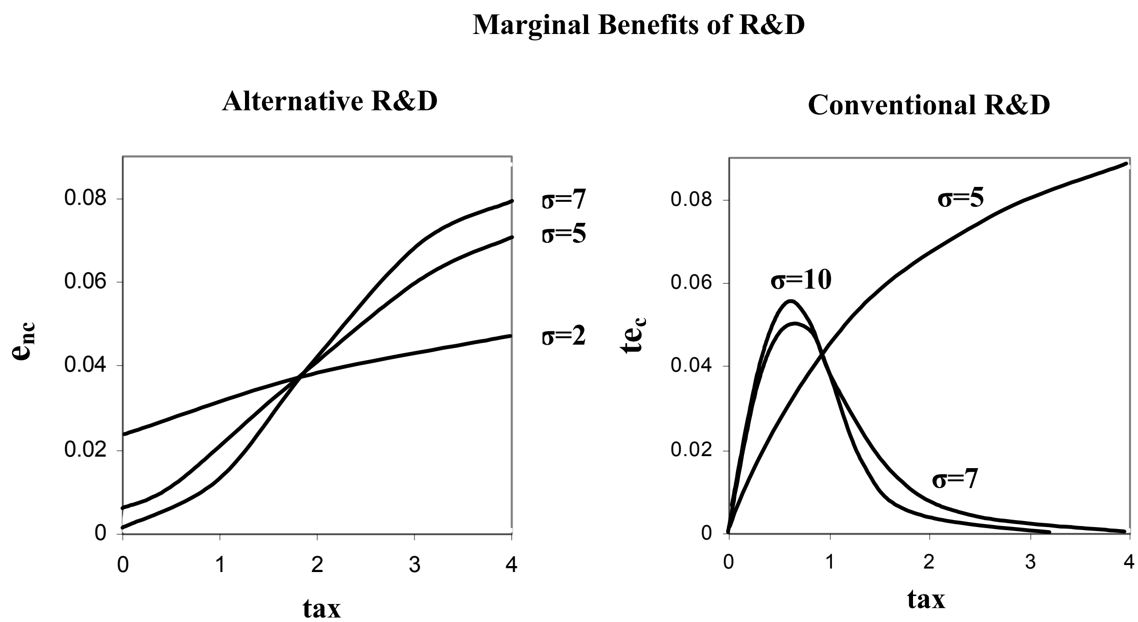

Figure 3: The left panel shows unconditional demand for non-carbon energy as a function of the carbon tax. The right panel shows the total tax expenditures as a function of the carbon tax. Price elasticity of demand is 1.1. Elasticity of substitution between energy and non-energy is .75.

an increase in R&D if and only if it increases the probability that the price of non-carbon energy will be less than the price of carbon energy. This is made precise in the following proposition.

Proposition 3 *Assume a firm has perfect substitution between carbon and non-carbon energy. Let α^* and $\tilde{\alpha}^*$ denote the optimal investment into R&D when the carbon tax has cumulative distribution function G and $\tilde{G}$, respectively. Then $\alpha^* \leq \tilde{\alpha}^*$ if and only if $G((1-\alpha^*)\eta) \geq \tilde{G}((1-\alpha^*)\eta)$.*

Proof. Under the assumption of perfect substitution, $e_{nc}^* = 0$ if $t < (1-\alpha)\eta$ and is equal to a constant, say M_α, if $t \geq (1-\alpha)\eta$ for any given level of α. If the CDF of t is G then $\mathrm{E}_t[e_{nc}^*] = \int_{(1-\alpha)\eta}^{\infty} M_\alpha dG = M_\alpha[1 - G((1-\alpha)\eta)]$. If $G((1-\alpha^*)\eta) \geq \tilde{G}((1-\alpha^*)\eta)$ then $g'(\alpha^*) \leq M_{\alpha^*}\left[1 - \tilde{G}((1-\alpha^*)\eta)\right]$, thus $\tilde{\alpha}^*$ must be larger than α^* to preserve optimally, and vice-versa. ■

The implication for firms is that, if they feel they are able to adjust their mix of energy inputs relatively quickly, then they should increase investments in ARD according to their beliefs about the probability that a carbon tax will make alternative energy competitive. The implication for public policy is that it may be beneficial to track the probability that the price of carbon will be such that alternative energy will become competitive.

5 Reducing Emissions in Conventional Technology

We return to investments in conventional technology R&D as described in Section 2.2. The right-hand side of the first order condition (2) simplifies to:

$$\mathrm{E}_t\left[\frac{\partial \pi}{\partial \alpha}\right] = -\,\mathrm{E}_t\left[t\frac{\partial \pi}{\partial p_c}\right] \tag{8}$$

Applying Hotelling's lemma, the first order condition (2) simplifies to

$$g'(\alpha) = \mathrm{E}_t\left[te_c^*\right] \tag{9}$$

where e_c^* represents the unconditional demand for carbon energy. The marginal benefit from CRD is equal to the total expected carbon tax expenditures of the firm. A firm will invest in R&D to reduce the emissions of its conventional technology to the degree that the firm expects total carbon tax expenditures to be high. Assuming an interior solution, if total tax expenditures increase in the carbon tax, then CRD increases in the expected value of the carbon tax. Similarly, if the expected tax expenditures increase in uncertainty, then optimal investment in CRD will increase in uncertainty. However, as we show in the next section, this problem is more complicated than ARD for two reasons: the marginal benefits to R&D are non-monotonic in the carbon tax, and corner-point solutions play a prominent role.

5.1 Increase in the expected value of a carbon tax: CRD

We find, similar to the result in Farzin and Kort, that CRD may decrease if the tax gets too high. Why? Because as the tax gets higher firms will substitute away from carbon energy, thus dampening the benefit from CRD. The marginal benefits to CRD (carbon tax expenditures) follow a Laffer curve – they first increase and then decrease in the carbon tax. In fact, it can be shown that for a monopolist with a CES production function and constant elasticity demand, there exists a t^* that maximizes carbon tax expenditures (see author for details). This effect in turn depends on the elasticity of substitution between carbon and other inputs and the elasticity of demand. The more elastic are substitution and demand, the sooner tax expenditures fall in the carbon tax. Thus, we would expect the *interior* solution for optimal CRD to first increase, and then decrease in a carbon tax.

We illustrate this using the same CES monopolist from Sections 4.1.1 and 4.2.1. The right panel of Figure 3 illustrates carbon tax expenditures at $\alpha = 0$, as a function of the carbon tax for three different values of σ. The point where tax equals 1 again represents the point where carbon and non-carbon energy become economically competitive. Again, the elasticity of substitution between energy and non-energy has little impact. The price elasticity of demand, however, has a significant impact on where the tax expenditures are maximized. For higher price elasticities and for higher elasticities of substitution, tax expenditures, and thus R&D, are maximized at a point that is before carbon and non-carbon energy become competitive.

Since the marginal benefits to CRD are non-monotonic, however, corner point solutions are important. First, if zero-emissions technology (i.e. $\alpha^* = 1$) is optimal for some $\hat{t} < \infty$, then it is optimal for every $t \geq \hat{t}$, even though the interior solution may fall. What this is saying is that, if 100% sequestration can be achieved for a reasonable cost, then optimal R&D will not fall even for a very large carbon tax.[7] Second, if zero-emissions technology is not optimal for some finite t, then as t gets large and the firm substitutes away from carbon energy, it is likely that the corner solution $\alpha^* = 0$ becomes optimal. Thus, firms will not gradually reduce investment in R&D as the expected tax increases, but rather will tend to drop quickly to zero investment once the expected tax gets too high.

We make these results explicit in the two general propositions below. In the first we assume that emissions from carbon energy cannot be reduced to zero (i.e. $g(1) = \infty$), and show that optimal CRD will decrease as t gets high. In the second, we show that if $g, g'$are finite at $\alpha = 1$, then if zero-emissions technology is not optimal for a finite t then optimal CRD will decrease as t gets high. Both of these results depend on the corner point solution where $\alpha^* = 0$. The logic is this: as long as α is strictly less than one, the cost of carbon energy will strictly increase in the carbon tax; for a high enough tax the firm will substitute almost entirely away from carbon energy (or stop producing altogether if they are not flexible enough to substitute). If a firm is not using carbon energy, then CRD has no value to them; thus the optimal amount of R&D at that point is zero. See Appendix for the proof of Proposition 4.

[7] This is important because AEP's official position is that 100% carbon sequestion is a realistic goal.

Proposition 4 *Assume that g is continuous, $g(0) = 0$ and $g(1) = \infty$, and that the total demand for carbon energy goes to zero as its price goes to infinity. Then optimal CRD will not monotonically increase in t.*

Proposition 5 *Assume that $g' < \infty$. Then either zero-emissions technology is optimal for a finite tax, or optimal R&D falls for a high enough carbon tax.*

Proof. Assume that $g(1) < \infty$. Proof by contradiction: assume that zero-emissions technology is never optimal for finite t but that α increases monotonically in t; according to the Karush-Kuhn-Tucker Theorem (KKT), that implies that for $\forall t \; g'(1) > te_c^*(t = 0)$, but this is impossible as long as $g' < \infty$. ■

The intuition behind both propositions is this: as long as $\alpha = 1$ is not optimal, then a high enough tax will drive the firm to stop using carbon energy altogether: the limit of profits as t goes to infinity is the same with and without R&D. This is not true for ARD: the limit of profits as t goes infinity is strictly different with and without R&D. For firms, optimal behavior depends on two things: beliefs about whether zero-emissions can be achieved, and beliefs about the magnitude of the carbon tax. If firms are flexible and believe that a carbon tax will be high enough to make non-carbon energy competitive with carbon energy, then R&D investments are better aimed at alternative rather than conventional technologies.

5.2 Mean-preserving spread in carbon tax: CRD

Since optimal CRD is low when the tax is low, and low when the tax is high, CRD will generally decrease in risk. The following proposition shows that there always exists some increase in risk that causes CRD to fall. Similar to ARD, optimal investment will fall when a low-probability, high-value tax is considered.

Proposition 6 *Optimal R&D into emissions reduction will decrease with some increases in risk regardless of the characteristics of the firm.*

Consider Figure 3 for further insight. For $\sigma = 7$,the marginal benefits of R&D under a certain tax $t = 1$ is about .05; the expected marginal benefits of R&D given a 50-50 chance of a carbon tax of 0 or 2 is an order of magnitude smaller, about .005. Note that if the expected carbon tax were quite high then optimal CRD would increase in risk: compare the marginal benefits of R&D under a certain tax of $t = 3$ with a 1/3 chance of $t = 1$ and a 2/3 chance of $t = 4$.

Combining the results from this section and Section 4, the implications for firms is that if they are flexible, an increase in risk in the neighborhood of the mean should induce more R&D into alternative technology, and less into conventional technology. If firms are not flexible, an increase in risk should induce less R&D across the board. For all firms, a low-probability, high-value tax should induce less R&D. Policy makers, and environmentalists, may want to keep in mind that an emphasis on risk may depress R&D investment, and in particular, an emphasis on a low-probability, high-value tax may cause inoptimally low R&D.

6 Energy Efficiency

In this section we briefly indicate how R&D into increased efficiency will fare under this framework. First we consider an improvement in the efficiency of carbon-energy, for example, an improvement in combined cycle gasification which will allow the firm to generate more energy per unit of carbon input. This is equivalent to reducing the effective price of e_c. The second period profit function would be $\pi\left(w,(1-\alpha)\left(p_c+t\right),p_{nc}\right)$ and the first order condition (2) simplifies to

$$g'(\alpha) = \mathrm{E}_t\left[\left(p_c+t\right)e_c^*\right] \tag{10}$$

The optimal response of this kind of R&D will be qualitatively similar to the optimal response of CRD: it will be non-monotonic in a carbon tax, and generally decrease in uncertainty. Note, however, that the threshold carbon tax, beyond which R&D decreases, will tend to be lower in this case, and strictly decreases in p_c. Note also, that the impact of risk in the price of non-carbon fuel will have the same impact as risk in a carbon tax. This implies that risky fuel prices should lead to less investment in fossil-fuel efficiency R&D.

Second we consider a general improvement in energy efficiency. If we assume that the production function for energy is homogenous of degree one, then energy efficiency would have the effect of reducing the effective price of both carbon and non-carbon energy. The second period profit function would be $\pi\left(w,(1-\alpha)\left(p_c+t\right),(1-\alpha)\,p_{nc}\right)$ and the first order condition (2) simplifies to

$$g'(\alpha) = \mathrm{E}_t\left[\left(p_c+t\right)e_c^*\right] + p_{nc}\,\mathrm{E}_t\left[e_{nc}^*\right] \tag{11}$$

Thus, the marginal benefit of R&D into energy efficiency is equal to the total expected cost of energy. The impact on energy efficiency R&D of an expected increase in a carbon tax will be qualitatively the same as the impact of an increase in the carbon tax on the overall cost of energy. This in turn depends on the characteristics of the firm in a similar manner as in the above analysis.

7 Conclusion

The first result is that if we consider R&D in terms of reducing the effective cost of inputs, then we can analyze the optimal amount of R&D in terms of how a change in price will change the unconditional demand for the inputs. An expected carbon tax will cause firms to invest in alternative R&D to the degree that a carbon tax will increase unconditional demand for non-carbon inputs. On the other hand, R&D into emissions-reduction of conventional technology will increase to the degree that a carbon tax raises total carbon tax expenditures. Similarly, an expected tax will increase R&D into energy efficiency to the degree that a carbon tax will increase the overall total cost of energy.

Taking this approach we have been able to show that environmental R&D will not monotonically increase in an expected carbon tax. If the elasticity of substitution between carbon and non-carbon energy is not high enough, then investment into R&D aimed at reducing the cost of alternative technologies will drop in a carbon tax. If, on the other hand, the elasticity of substitution is high enough, then this kind of R&D will increase monotonically in a carbon tax. If zero-emissions technology is not optimal, then R&D into conventional technologies responds non-monotonically to a carbon tax, first increasing, then decreasing as the tax gets too high: there is a threshold tax above which optimal R&D decreases. The level of this threshold tax decreases as both demand and production get more elastic.

We go on to analyze the firm's optimal response to an increase in risk in the carbon tax. Firms that have flexibility to substitute between carbon and non-carbon energy in a relatively short time should increase investment in alternative R&D with increases in risk that increase the probability that non-carbon energy will be competitive with carbon energy. For firms that have less flexibility to substitute, R&D into alternative energy decreases in risk. Similarly, the optimal level of conventional R&D will generally decrease in risk. Finally, for all firms and for both R&D programs, a low-probability, high value tax induces a reduction in the optimal R&D.

Our results emphasize the fact that different R&D programs have different optimal responses to policy. Thus, when modeling induced technical change, care must be taken. In particular, technological change may have complex impacts on the marginal cost of abatement, depending on firm characteristics and on the R&D program.

For policy makers and environmentalists, the implication is that balanced information about climate change will induce the most R&D: extreme events should not be over-emphasized. In future work, we will study the gap between profit-maximizing R&D and socially optimal R&D to make recommendations about how optimal technology policy is impacted by risk.

For our examples we have used a monopolist with CES production and demand, but Propositions 2, 4, 5, and 6 are valid for any production function and for a competitive equilibrium as well as for a monopoly. We have made a number of simplifying assumptions. We have used a framework with fixed prices for the inputs. The results remain unchanged, however, if we allow prices to increase with demand for the input. We consider uncertainty about the magnitude of a tax, but not the timing. This can be partially justified by the idea that R&D projects can take a long time; but future work on the timing of R&D investments is needed. We assume that the returns from R&D are deterministic. It will be important to include both uncertainty in policy and uncertainty in returns to R&D to determine how they interact. We have assumed risk neutrality. Initial theoretical work indicates that firms will be more likely to increase R&D spending in response to an increase in risk as they get more prudent; however, our numerical examples showed no impact of risk aversion even for very high levels of prudence. Another key issue that must be addressed is that the providers of R&D may be different from those who implement the new technology. If the producers of R&D hold significant market power, it may impact the results. We have used a price-based mechanism to represent policy, and considered a firm's response to a MPS in this price. An alternative is to consider an MPS in an emissions standard. We note, however, that if ex-post policy is efficient, and both the tax and the standard are based on underlying scientific uncertainty, the results are qualitatively the same.

A Appendix

A.1 Proof of Proposition 1

Proof. Define $\zeta \equiv \frac{1}{1-\rho}$ to be the elasticity of substitution between energy and non-energy inputs and $\sigma \equiv \frac{1}{1-\gamma}$ as the elasticity of substitution between carbon and non-carbon energy inputs. To simplify the notation let the total price of carbon inputs be represented by $\tilde{p}_c$ and the total price of non-carbon inputs be represented by p_{nc}. Then the unconditional demand for non-carbon inputs is

$$e_{nc}^{*} = p_{nc}^{\frac{1}{\gamma-1}} \left[\tilde{p}_c^{\frac{\gamma}{\gamma-1}} + p_{nc}^{\frac{\gamma}{\gamma-1}} \right]^{\frac{\gamma-\rho}{(\rho-1)\gamma}} \left(w^{\frac{\rho}{\rho-1}} + \left[\tilde{p}_c^{\frac{\gamma}{\gamma-1}} + p_{nc}^{\frac{\gamma}{\gamma-1}} \right]^{\frac{\rho}{\rho-1} \frac{\gamma-1}{\gamma}} \right)^{\frac{-1-\rho b+b}{\rho}} \left(\frac{b}{b-1} \frac{1}{A} \right)^{-b} \tag{12}$$

The sign of the derivative of (12) with respect to $\tilde{p}_c$, is equal to the sign of:

$$\frac{\gamma}{\gamma-1} \left\{ \frac{\gamma-\rho}{(\rho-1)\gamma} w^{\frac{\rho}{\rho-1}} + \left(\frac{\gamma-\rho}{(\rho-1)\gamma} + \frac{-1-\rho b+b}{\rho} \frac{\rho}{\rho-1} \frac{\gamma-1}{\gamma} \right) \left[\tilde{p}_c^{\frac{\gamma}{\gamma-1}} + p_{nc}^{\frac{\gamma}{\gamma-1}} \right]^{\frac{\rho}{\rho-1} \frac{\gamma-1}{\gamma}} \right\} \tag{13}$$

This quantity can be unambiguously signed for all values of w, p_c and p_{nc} only when the coefficients of both terms inside the curly brackets have the same sign. Consider $\gamma > 0$. Then $\frac{\gamma}{\gamma-1} < 0$ (since $\gamma < 1$). $\frac{\gamma-\rho}{(\rho-1)\gamma} < 0$ if $\gamma > \rho$. The second coefficient inside the curly brackets simplifies to

$$= \frac{-b\gamma + b - 1}{\gamma} \tag{14}$$

The quantity in (14) is negative if

$$\gamma > \frac{b-1}{b} \tag{15}$$

This calculation can be repeated for the other 3 cases. Thus, the derivative of (12) with respect to $\tilde{p}_c$ (and therefore with respect to t) is positive for all values of w, p_c, and p_{nc} if $\gamma > \overline{\gamma}$ where

$$\overline{\gamma} = \max\left[\rho, \frac{b-1}{b}\right] \tag{16}$$

and is negative if $\gamma < \underline{\gamma}$ where

$$\underline{\gamma} = \min\left[\rho, \frac{b-1}{b}\right] \tag{17}$$

If $\underline{\gamma} < \gamma < \overline{\gamma}$ then the sign of the derivative depends on the relative values of the prices. In terms of the elasticities of substitution, this says that $\sigma \geq \max[\zeta, b]$ is a sufficient condition for unconditional demand for non-carbon energy to increase in t.

Moreover we show that if there exists a $\hat{t}$ for which the interior solution is optimal (as opposed to the corner point solution $\alpha^* = 0$), then the interior solution is optimal for all $t \geq \hat{t}$.Thus, if the interior solution is increasing in t, the global solution is increasing in t. Define $\alpha(t)$ as the interior solution to (1). If $\alpha^* = 0$ is the global maximum at tax t, then

$$-g(\alpha(t)) + \pi(w, p_c + t, p_c + (1 - \alpha(t))\eta) < \pi(w, p_c + t, p_c + \eta) \tag{18}$$

But, under the assumptions of the proposition, the right-hand side of (18) is decreasing in t faster than the left-hand side[8], thus if the interior solution $\alpha(\hat{t})$ is the global maximum for $t = \hat{t}$, then the interior solution $\alpha(t)$ is the global maximum for $t \geq \hat{t}$. ∎

It is interesting to note what happens if neither condition (16) nor (17) holds. If $\zeta < \sigma < b$, then if p_c is high enough[9], R&D will decrease monotonically in the carbon tax; if p_c is low, R&D will first increase and then decrease. On the other hand, if $b < \sigma < \zeta$, then if p_c is high enough, R&D will increase monotonically; if p_c is low then R&D will first decrease and then increase.

A.2 Proof of Proposition 2

Proof. Optimal R&D will increase unambiguously in risk if and only if e^*_{nc} is everywhere convex. The assumption that R&D increases in the carbon tax implies that e^*_{nc} increases in t. On the other hand, e^*_{nc} must be bounded from above, and thus cannot be everywhere convex. ∎

A.3 Proof of Proposition 4

Proof. Define $\alpha(t)$ as a local maximum (i.e. it satisfies the first and second order conditions) of the decision problem (1) assuming CRD, and assume that it is monotonically increasing in t. If

[8] The derivative of both sides with respect to t is $-e^*_c$. Under the assumptions of the proposition, e^*_c is increasing in $p_{nc} = p_c + (1 - \alpha)\eta$. p_{nc} is strictly higher when $\alpha = 0$, so the slope of the rhs is more negative.

[9] Specifically, if p_c satisfies $\left|\frac{-b\gamma+b-1}{\gamma}\right| \left[p_c^{\frac{\gamma}{\gamma-1}} + p_{nc}^{\frac{\gamma}{\gamma-1}}\right]^{\frac{\rho}{\rho-1}\frac{\gamma-1}{\gamma}} > \left|\frac{\gamma-\rho}{(\rho-1)\gamma}\right| w^{\frac{\rho}{\rho-1}}$

these assumptions do not hold then Proposition 4 is true by default. Since $\alpha(t)$ is monotonically increasing in t and is bounded above by 1, the limit of α exists. First, assume that $\lim_{t\to\infty}\alpha(t) = 1$, then consider

$$\lim_{t\to\infty} -g(\alpha(t)) + \pi((1-\alpha(t))t) \tag{19}$$

where we are suppressing the prices p_c, p_{nc}, and w in the profit function. Note that $0 \leq \pi((1-\alpha(t))t) \leq \pi(0)$: profits are bounded below by zero and above by business-as-usual profits. Since g is continuous, $\lim_{t\to\infty} -g(\alpha(t)) = -\infty$, so (19) goes to negative infinity as t gets large; since $\pi(t) \geq 0$, the corner point solution $\alpha^* = 0$ is optimal for large enough t.

Now assume that $\lim_{t\to\infty}\alpha(t) = c < 1$. The profit function is continuous [25], thus

$$\begin{aligned} & \lim_{t\to\infty} -g(\alpha(t)) + \pi((1-\alpha(t))t) && (20)\\ = \;& -g(c) + \pi((1-c)\infty) && (21)\\ < \;& -g(0) + \pi(\infty) = \pi(\infty) && (22)\end{aligned}$$

Note that (22) is the value of the maximand at the corner point solution $\alpha^* = 0$. ■

A.4 Proof of Proposition 6

Proof. In the uninteresting cases where $\alpha^* = 0$ or 1 for all t, then every increase in risk is weakly decreasing α^*. Assume that $\alpha^* = 0$ when $t = 0$ and there exists a $\hat{t}$ such that optimal CRD $\hat{\alpha} > 0$.

For every t^H there exists a $p = \frac{\hat{t}}{t^H}$ such that $\hat{t} = p * t^H + (1-p) * 0$. Consider

$$\begin{aligned} \alpha^* &\equiv \arg\max_{\alpha} -g(\alpha) + p\pi\left((1-\alpha)t^H\right) + (1-p)\pi(0) && (23) \\ &= \arg\max_{\alpha} -g(\alpha) + p\pi\left((1-\alpha)t^H\right) && (24) \\ &= \arg\max_{\alpha} -g(\alpha) + \frac{\hat{t}}{t^H}\pi\left((1-\alpha)t^H\right) && (25) \end{aligned}$$

Since profits are strictly bounded,

$$\lim_{t^H \to \infty} -g(\alpha) + \frac{\hat{t}}{t^H}\pi\left((1-\alpha)t^H\right) = -g(\alpha) \tag{26}$$

which is strictly negative for $\alpha > \hat{\alpha}$. Thus the optimal α^* must decrease for t^H high enough, i.e. we have shown that there exists a MPS around $\hat{t}$ that induces a decrease in α^*. ■

References

[1] Susan Athey and Armin Schmulzer. Product and process flexibility in an innovative environment. *Rand Journal of Economics*, 26:557–574, 1995.

[2] Erin Baker. Uncertainty and learning in a strategic environment: Global climate change. *Resource and Energy Economics*, Forthcoming.

[3] P. Buonanno, C. Carraro, and M. Galeotti. Endogenous induced technical change and the costs of kyoto. *Resource and Energy Economics*, 25:11–34, 2003.

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

[5] Y. H. Farzin and P.M. Kort. Pollution abatement investment when environmental regulation is uncertain. *Journal of Public Economic Theory*, 2:183–212, 2000.

[6] Carolyn Fischer, Ian W.H. Parry, and William A. Pizer. Instrument choice for environmental protection when technological innovation is endogenous. *Journal of Environmental Economics and Management*, 45:523–545, 2003.

[7] Don Fullerton, Shaun P. McDermott, and Jonathan P. Caulkins. Sulpher dioxide compliance of a regulated utility. *Journal of Environmental Economics and Management*, 24:32–53, 1997.

[8] Christian Gollier, Bruno Jullien, and Nicolas Treich. Scientific progress and irreversibility: An economic interpretation of the 'precautionary pricipal'. *Journal of Public Economics*, 75:229–253, 2000.

[9] Larry H. Goulder, Ian Parry, R.C. Williams, and D. Butraw. The cost-effectiveness of alternative instruments for environmental protection in a second-best setting. *Journal of Public Economics*, 72:329–360, 1999.

[10] Lawrence Goulder and Koshy Mathai. Optimal CO2 abatement in the presence of induced technical change. In Arnulf Grubler, Nebojsa Nakicenovic, and William D. Nordhaus, editors, *Technological Change and the Environment*, pages 210–250. RFF and IIASA, Washington D.C and Laxenburg, Austria, 2002.

[11] Chulho Jung, Kerry Krutilla, and Roy Boyd. Incentives for advanced pollution abatement technology at the industry level: An evaluation of policy alternatives. *Journal of environmental economics and management*, 30:95–111, 1996.

[12] K. Keller, B. Bolker, and D. Bradford. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management*, Forthcoming, 2004.

[13] Charles D. Kolstad. Learning and stock effects in environmental regulation: The case of greenhouse gas emissions. *Journal of Environmental Economics and Management*, 31:1–18, 1996.

[14] Bruce Larson. Environmental policy uncertainty and investment in pollution control. University of Connecticut, Department of Agricultural and Resource Economics, 2004.

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

[16] Juan-Pablo Montero. Permits, standard, and technology innovation. *Journal of Environmental Economics and Management*, 44:23–44, 2002.

[17] William D. Nordhaus. Modeling induced innovation in climate change policy. In Arnulf Grubler, Nebojsa Nakicenovic, and William D. Nordhaus, editors, *Technological Change and the Environment*, pages 182–209. RFF and IIASA, Washington D.C. and Laxenburg, Austria, 2002.

[18] Ian Parry. Pollution regulation and the efficiency gains from technological innovation. *Journal of Regulatory Economics*, 14:229–254, 1998.

[19] David Popp. ENTICE-BR: The effects of backstop technology R&D on climate policy models. Technical Report 10285, NBER, 2004.

[20] David Popp. ENTICE: Endognenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management*, Forthcoming.

[21] Michael Rothschild and Joseph Stiglitz. Increasing risk I: A definition. *Journal of Economic Theory*, 2:225–243, 1970.

[22] Michael Rothschild and Joseph Stiglitz. Increasing risk II: Its economic consequences. *Journal of Economic Theory*, 3:66–84, 1971.

[23] Ian Sue Wing. Induced technical change and the cost of climate policy. Technical Report 102, MIT Joint Program on the Science and Policy of Global Change, 2003.

[24] Alistair Ulph and David Ulph. Global warming, irreversibility and learning. *The Economic Journal*, 107:636–650, 1997.

[25] Hal R. Varian. *Microeconomic Analysis*. W. W. Norton & Company, New York, N.Y., 1992.

[26] M. Webster. The curious role of "learning" in climate policy: Should we wait for more data? *Energy Journal*, 23(2):97–119, 2002.