Optimal Technology R&D in the Face of Climate Uncertainty

Erin Baker

Asst. Professor

University of Massachusetts

220 ELab

Amherst, MA 01002

edbaker@ecs.umass.edu

Leon Clarke

Joint Global Change Research Institute

John Weyant

Professor

Stanford University

January 14, 2004

Abstract

This paper explores optimal near-term technology R&D in the face of uncertain damages caused by the buildup of greenhouse gases. The paper puts particular emphasis on understanding how optimal near-term R&D expenditures might vary based on the technologies pursued in the R&D program. The exploration is conducted in the context of varying impacts from R&D on the global abatement cost function. The R&D planning problem is considered first within a theoretical framework and is then pursued in a stylized application using the DICE model. The paper provides intuition into the circumstances under which near-term technology R&D might increase or decrease under uncertainty, thereby serving as a hedge against climate uncertainty.

1 Introduction

Climate change presents a challenge rife with uncertainty. Near-term decisions must be made in the face of, among others, uncertainties in population growth, economic growth, technological advance, and the ultimate human welfare impacts of increasing concentrations of greenhouse gases (GHGs). This paper explores the response of one important near-term decision variable—optimal near-term global R&D expenditures—to increasing uncertainty in one critical dimension—the human welfare impacts of increasing GHG concentrations, which we term "climate uncertainty" for convenience.

The paper is motivated in large part by the argument that R&D into technologies with zero or very-low emissions might be viewed as a hedge against uncertainty. The more uncertain we are about the potential consequences of increasing GHG concentrations, the more we may want to invest in improving technologies that would reduce the costs of dramatic emissions reductions. If the future resolution of uncertainty indicates that quick action is warranted—for example, if climate damages will be very severe unless the stock of emissions is stabilized at a low ceiling, say 450 or lower—having these technologies available at lower cost will allow for lower cost action. Hence, R&D may provide an adaptive approach to climate policy by giving us more flexibility to act in the future after we learn.

The R&D planning problem is complicated by the fact that there are many possible R&D programs. Many technologies might be pursued and in varying proportions, from photovoltaic cells to technologies for separating carbon dioxide from power plant flue gases to improvements in the efficiency of automobiles and electric power plants. A research program to improve the efficiency of coal-fired electricity generation will create a very different abatement cost structure than an R&D program into photovoltaic cells. The first program will lower the cost for moderate reductions in GHG emissions; the second program will lower the costs of severely reducing emissions. This raises the possibility that the response of R&D to increasing uncertainty may not be independent of the R&D program —it may depend on what sort of R&D we are considering. Hence, an important objective of this paper is to understand how the choice of R&D program impacts the response to increasing uncertainty.

The paper proceeds first with a general discussion of the manner in which R&D might alter the global abatement cost function. Three illustrative examples are proposed for further consideration. Next, a two-period theoretical model is used to better understand whether and when optimal R&D might increase or decrease in increasing uncertainty. Last, the three illustrative examples are pursued in a stylized, empirical application using a stochastic version of the DICE model.

The primary theme of the analysis is that there is no unambiguous answer to the central question it addresses. Put simply, optimal R&D might increase or decrease with uncertainty depending on the nature of existing technology, the manner in which R&D alters technology, and way in which uncertainty increases. In support of this, the theoretical analysis indicates that no R&D program will serve as a hedge against all forms of increasing uncertainty, and the stylized empirical analysis showed both increasing and decreasing optimal R&D. Taken in total, however, the analysis does indicate that R&D into very-low emissions technologies can be a hedge against some kinds of uncertainty. In some sense, the issue is less a matter of increasing or decreasing uncertainty, and more a matter of whether a change in uncertainty changes the distribution of probability mass toward or away from the technologies affected by a particular R&D program. Optimal R&D into very-low emissions technologies will increase with changed assessments of an uncertain future to the extent that these changed perceptions put a greater probability mass on the need to use these technologies, that is, on very high levels of abatement.

The remainder of this paper proceeds as follows. Section 2 provides relevant background. Section 3 discusses the manner in which R&D might impact the abatement cost function and introduces the three example R&D programs. Section 4 provides the theoretical discussion of optimal global R&D response to climate uncertainty in a two-period model. Section 5 introduces the stylized, empirical application and Section 6 discusses the results of the analysis. Section 7 sums up.

2 Background

Two lines of research are appropriate background for this paper. First, the paper is parallel to the line of research that considers the impacts of uncertainty and learning on optimal near-term abatement levels. This work considers the impacts of learning given uncertainty, the impacts of increasing uncertainty without learning, and the impacts of increasing uncertainty given that we will learn.[1] In general, this line of work indicates lower near-term action when there is uncertainty with the possibility of learning. However, the results are far from absolute. For example, both Webster (2002) and Gollier et al. (2000) show that the results can be reversed under certain condition. The work in this paper differs from the above work by its consideration of near-term R&D rather than near-term abatement. However, to the extent that R&D investments can be viewed in a similar light as near-term abatement investments—both are costly in the near-term and improve our position if damages turn out to be higher than expected—this paper follows closely from this literature.

A second line of relevant research is the growing body of work on endogenous technological advance in the context of climate change. This literature covers technological change that is in some way induced by policy, generally by the indirect effect on market actors, but also as a control variable.[2] Learning-by-doing is by far the most prominent of the approaches, but several authors have also considered endogenous R&D. Two important thrusts of this work are to understand the implications of endogenous technological change on the costs of abatement and on the timing of abatement, and therefore the stringency of near-term action. The work in this paper differs from the majority of this previous work in that we consider R&D as a control variable (as opposed to output from an imperfectly behaving

[1] A non-exhaustive list of this literature includes Baker (2003), Gollier (2000), Keller, Bolker, and Bradford (2004), Kolstad (1996), Manne (1996), Pizer (1999), Ulph & Ulph (1997), and Webster (2002).

[2] On the modeling front, this literature includes, among others, Buonanno, Carraro, and Galleotti (2003), Goulder and Mathai (2000), Goulder and Schneider (1999), Manne and Richels (2002), Nordhaus (2002), Popp (2002), Schneider and Goulder (1997), Wing (2003), and van der Zwaan et al. (2002). On an empirical front, this literature includes Newell (1997), and Newell, Jaffe, and Stavins (1998). Clarke and Weyant (2002), and Jaffe, Newell, and Stavins (2001) both provide useful surveys of the literature.

market), and that we consider optimal R&D under uncertainty. This paper therefore follows most closely the efforts of those authors that have considered R&D as a control variable, including Goulder and Mathai (2000), Nordhaus (2002), and Popp (2002).[3]

3 Representing the Impacts of R&D

This section discusses the manner in which R&D impacts technology through the lens of the simplifying construct of a global abatement cost function. As we are ultimately concerned in this paper with the socially optimal *global* level of R&D, our concern lies with impacts on the *global* abatement cost function, as opposed to firm or country-level abatement cost functions. The section begins with a general discussion followed by three illustrative examples that will be touched on in the theoretical discussion in Section 4 and will serve as the basis for the stylized empirical work in sections 5 and 6.

To better understand the impact of R&D, we begin by tying the global abatement cost function to the global production function. The production function often provides a clearer intuition into the real-world technologies of which it is an abstraction than does the abatement cost function, and this is valuable when thinking about R&D. Figure 1 shows a representative isoquant of a stylized global production function, along with the corresponding abatement cost function. The production function underlying the isoquant is based on two inputs: "standard" inputs, whose costs are internal to the market absent any environmental policy, represented by τ; and environmental inputs, the costs of which would not, without government intervention, make their way into private decisions, represented by ε. As an example, electricity generated from coal has standard inputs of fuel, labor, maintenance, and infrastructure, and a set of environmental inputs associated with carbon emissions, sulfur emissions, particulate emissions, and so forth. Hence, the production function can be written as $Q = f(\tau, \varepsilon)$. As

[3] Note that while Popp and Nordhaus include energy R&D as a control variable, they also assume that the costs of energy R&D are supernormal so as to capture the "crowding out" of R&D in other sectors where returns are supernormal.

this paper is concerned with climate change, ε is henceforth assumed to represent GHG emissions.

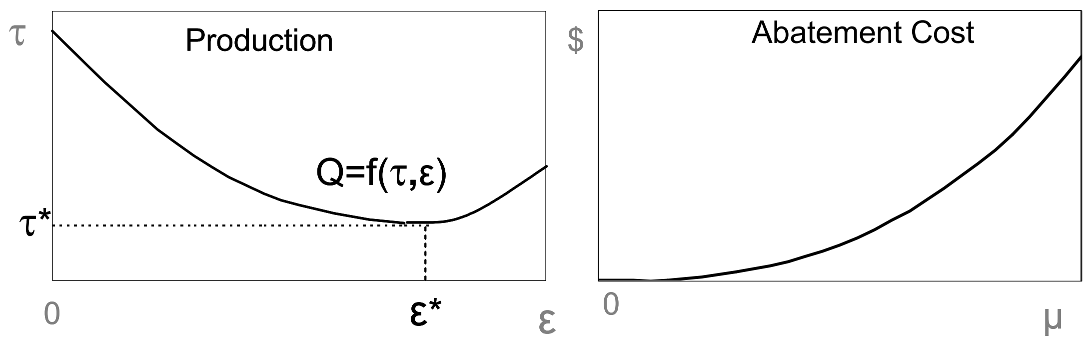

Figure 1: The left hand-panel shows a representative isoquant. τ and ε represent standard and environmental inputs, respectively. The right-hand panel shows the associated abatement cost function. μ is the level of emissions abatement.

As Figure 1 shows, for any particular production level Q, there exists an emissions level, ε^*, at which standard inputs are minimized. If emissions are increased above ε^*, standard input requirements increase; in other words, it costs more to emit more. The minimum standard input requirements are given by $\tau^* = \tau(\varepsilon^*)$, where $\tau(\cdot)$ is the implicit function derived from setting the Q equal to a constant. In the absence of emissions policies, τ^* represents the predominant mix of technologies in the economy.[4]

The abatement cost function represents the change in standard input requirements associated with a reduction in emissions below ε^*. Emissions reductions might be achieved through two mechanisms: (1) by reducing output and maintaining a constant ratio of standard inputs to emissions, and (2) by substituting standard inputs for emissions while maintaining constant output. In this paper, the latter paradigm is assumed: the abatement cost function corresponds to the cost of reducing emissions through substitution.[5] Without loss of generality, let $\varepsilon^* = 1$. Let $c(\mu)$ be the abatement cost curve associated with a particular production function, where $\mu = 1 - \varepsilon$ is the fractional emissions reduction. Hence,

[4] Note that this point will include both emitting and non-emitting technologies. For example, photovoltaic cells and wind power are both economic in particular applications absent emissions policies. Nevertheless, we expect that the majority of technologies will be those that emit GHGs.

[5] Functionally, abatement will reduce global production because, with a fixed stock of standard inputs, any additional inputs must come from output.

$$c(\mu) = \tau(1-\mu) - \tau^* \tag{1}$$

One additional characteristic of the isoquant in Figure 1 bears mentioning. For the economy represented in the figure, there exists a GHG-free alternative technology with finite cost, so the isoquant intersects the vertical axis and full abatement can be obtained at finite cost. The point of intersection represents the lowest cost GHG-free technology available, such as wind-power or solar power.[6] This notion of a GHG-free technology with finite cost is consistent with the DICE model, which will be explored in Sections 5 and 6.

This paper is concerned with R&D. Technological advance through successful R&D might shift, stretch, and deform the production function, and therefore the abatement cost function, in an infinite number of ways. The remainder of this section introduces three illustrative examples of R&D effects and suggests possible real-world analogs for each. Note that the three examples are by no means an exhaustive set.

3.1 Program 1: Constant Emissions Reduction

The left panel in Figure 2 illustrates a leftward shift in the production function by an amount α. This representation of R&D is found in Baudry (2000) as a shift in the cost function. It allows a given level of output to be produced using the same standard input levels as in the original production function, but with emissions lowered by the absolute amount α.

Let quantities associated with the new production function be marked with a $\tilde{}$. Then the new production function can be related to the old production function by noting that $\tilde{\tau}(\varepsilon) = \tau(\varepsilon + \alpha)$. Let $\tilde{c}(\mu; \alpha)$ represent the cost function after R&D. Thus the new cost function, shown in the right panel of

[6] In reality, there are GHG emissions associated with virtually every technology. For example, the construction of wind-turbines requires energy, a proportion of which will have come from freely-emitting fossil fuels. The notion of a perfectly GHG-free technology is a benchmark for analysis.

Figure 2, is given by

$$\tilde{c}(\mu;\alpha) = \tau(\varepsilon+\alpha) - \tau^* = c(1-\varepsilon-\alpha) = c(\mu-\alpha) \tag{2}$$

It is productive to consider how this abstraction might roughly correspond to real-world R&D programs. One interpretation is that the R&D program results in the development of a set of technologies with identical standard input requirements of those associated with τ^* but with zero emissions and of only a limited total capacity. And these technologies must be *in addition* to the stock of technologies already available. So, for example, the development of a limited capacity of tidal power at a cost comparable to that of coal-fired power plants today would fit this paradigm. Another interpretation might be a no-cost sequestration program.

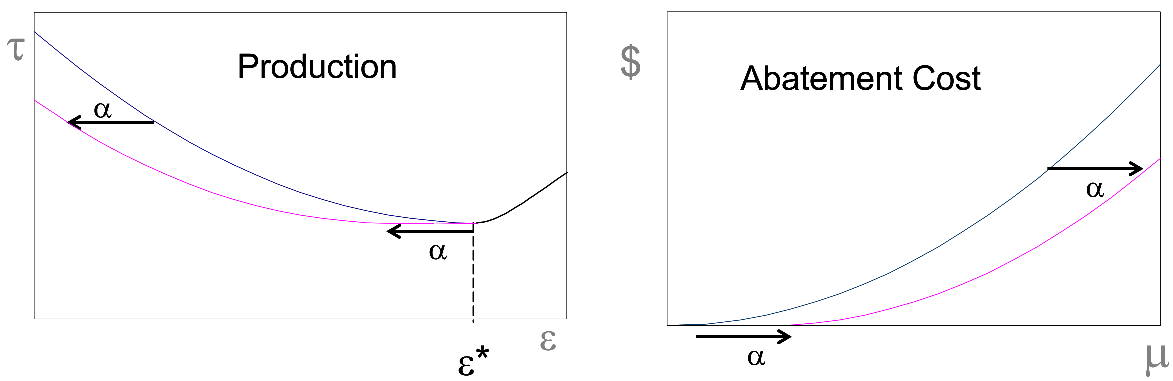

Figure 2: Program 1: Constant Emissions Reduction. The left-hand panel shows a representative isoquant before and after R&D. The right-hand panel shows the associated abatement cost function.

3.2 Program 2: Cost Reduction

Figure 3 illustrates an R&D program that shifts the production function downward toward τ^* by a fixed percentage, again denoted by α. This representation is found in Goulder and Mathai (2002). It allows for proportionally lower costs of reducing emissions for all emissions levels.

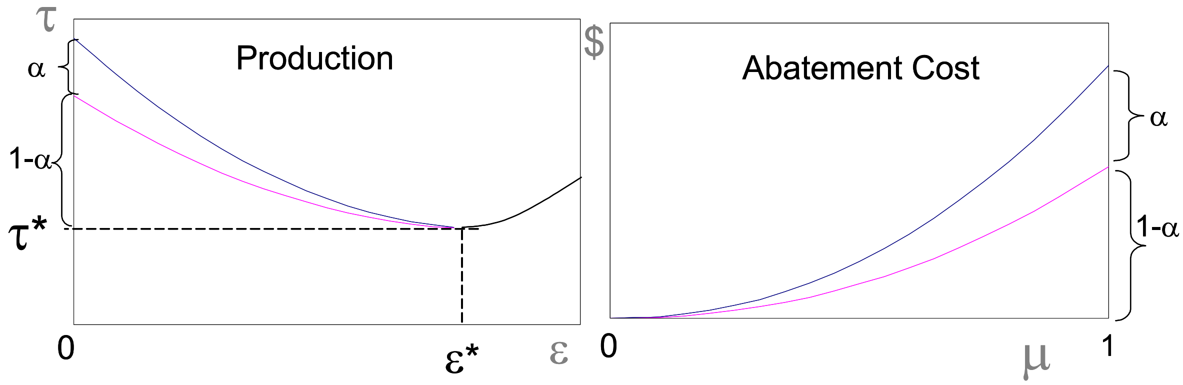

Figure 3: Program 2: Cost Reduction. The left-hand panel shows a representative isoquant before and after R&D. The right-hand panel shows the associated abatement cost function.

The new production relationship is given by $\tilde{\tau}(\varepsilon) = \alpha\tau^* + (1-\alpha)\tau(\varepsilon)$. Thus the new cost function, shown in the right panel of Figure 3, is given by

$$\tilde{c}(\mu;\alpha) = \alpha\tau^* + (1-\alpha)\tau(\varepsilon) - \tau^* = (1-\alpha)c(\mu) \tag{3}$$

In general, this program should be interpreted as primarily reducing the cost of the more radical, alternative energy sources. One interpretation is that it aims to, for example, halve the cost differential between all of the lower-carbon options and the current option. In this interpretation, R&D is spread among a multitude of technologies. Alternatively, it can be interpreted as approximating a reduction in the cost of no-carbon alternatives. The reduction is manifest across the full isoquant because the no-carbon alternatives are used across the isoquant, but in decreasing intensity. For example, a reduction in the cost of photovoltaic cells would reduce the cost of reaching zero emissions, but would also reduce the costs of achieving small decreases in emissions because they could be more cheaply deployed in smaller, niche markets.

3.3 Program 3: Proportional Emissions Reduction

Figure 4 illustrates an R&D program that decreases the emissions-output ratio by a proportional amount α. For any given level of output and standard inputs, R&D reduces the emissions (by a fixed percentage) needed to achieve that output. Thus, the production function is shifted leftward by a constant percentage, as opposed to a constant amount in Program 1. This can be viewed as the inverse of Program 2 in the sense that the R&D has the largest effect on the higher emitting portions of the production function. It has no effect on the costs of moving all the way to a zero-emissions world, and therefore on zero-carbon technologies. This representation of R&D is found in a number of papers on technical change (e.g. Farzin and Kort, 2000), including Buonanno, Carraro, and Galeotti (2003) and Nordhaus (2002) as a reduction in the emission-output parameter.

The new production relationship is given by $\tilde{\tau}(\varepsilon) = \tau(\frac{\varepsilon}{1-\alpha})$. Thus the new cost function, shown in the right panel of Figure 4, is given by

$$\tilde{c}(\mu, \alpha) = \tau(\frac{\varepsilon}{1-\alpha}) - \tau^* = c(1 - \frac{\varepsilon}{1-\alpha}) = c(\frac{\mu - \alpha}{1-\alpha}) \tag{4}$$

It is interesting to note that this program does not reduce marginal cost everywhere across the abatement cost function, whereas it is commonly assumed that environmental R&D will decrease marginal costs everywhere. In this case, R&D significantly reduces the cost of small amounts of abatement, while leaving the costs of large amounts of abatement virtually unchanged. Hence, the marginal cost of moving from small to large amounts will be increased.

In general, this program can be interpreted as primarily reducing emissions in the currently economic alternatives, such as combined cycle gasification for electric generation. A second interpretation is that it represents a multi-technology funding approach intended to decrease the emissions-output ratio of all technologies. The impacts would therefore be larger for technologies with larger emissions, and the program would include no funding for zero-emissions technologies.

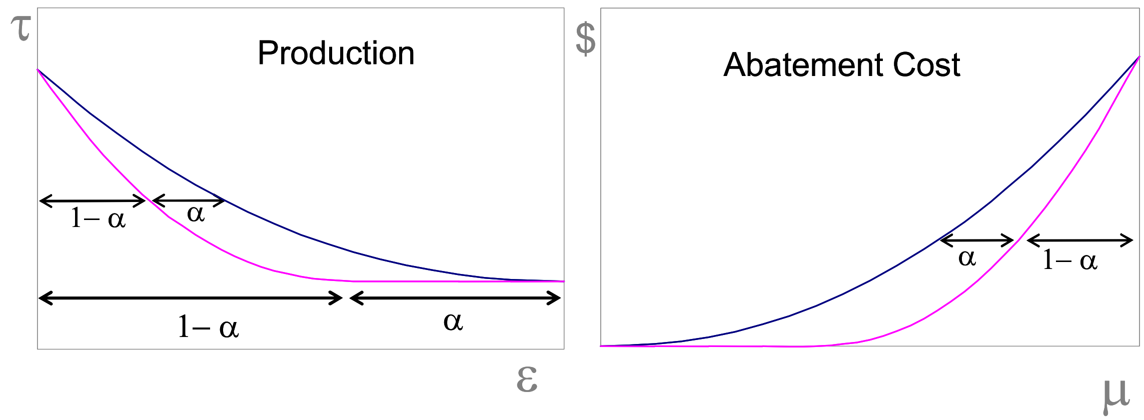

Figure 4: Program 3: Proportional Emissions Reduction. The left-hand panel shows a representative isoquant before and after R&D. The right-hand panel shows the associated abatement cost function.

4 Theory

4.1 The Model

This section uses a two-period, social planning model to explore the theory underlying the response of R&D to increasing uncertainty in climate damages. As we will show, the analysis indicates that, at a theoretical level, the impacts on optimal R&D will be ambiguous except in very specific cases; in other words, the parameters will matter.

The model proceeds as follows. R&D investments are made in the first period. The benefits of R&D, in terms of alterations to the abatement cost function, accrue in the second period. At the time of the R&D investment, the extent of damages resulting from emissions in the second period, and hence the benefits of alterations to the abatement cost function through R&D, are uncertain. This uncertainty is resolved at the beginning of the second period, allowing for perfectly informed emissions abatement from that point forward using the altered abatement cost function. For simplicity, it is assumed that there are no emissions in the first period.

Consistent with Section 3, α represents the effects of R&D on the production function, and therefore the abatement cost function. The (opportunity) costs of achieving a particular level of advance are given by $g(\alpha)$, which is assumed to be increasing and convex: $\frac{\partial g}{\partial \alpha} > 0$ and $\frac{\partial^2 g}{\partial \alpha^2} > 0$.

The objective of the social planning model is to minimize the sum of the costs of R&D, the expected global costs of abatement, and the expected global damages from emissions. The social planning problem is therefore given by

$$\min_{\alpha} \quad g(\alpha) + \mathrm{E}_z \left\{ \min_{\mu} \quad c(\mu;\alpha) + D(\mu;z) \right\} \tag{5}$$

where $D(\mu;z)$ are the climate damages in the second period resulting from the total stock of emissions, z is a stochastic parameter meant to capture uncertainty in climate damages, and E_z is the expectation operator over z. Let $V(\alpha,z) = c(\mu^*;\alpha) + D(\mu^*;z)$ be the second-period welfare impact based on optimal abatement, μ^*, for a given level of technological change, α, and resolution of climate uncertainty, z. The social planning problem can now be written as.

$$\min_{\alpha} \quad g(\alpha) + E_z\{V(\alpha,z)\} \tag{6}$$

Now define an increase in uncertainty as a mean-preserving-spread in the spirit of Rothschild and Stiglitz (1970).[7] Whether optimal R&D, and therefore optimal technological advance, α^*, increases or decreases with increasing uncertainty depends on the convexity of $\frac{\partial V}{\partial \alpha}$ with respect to the uncertain damage parameter, z (Rothschild and Stiglitz, 1971). α^* will increase (decrease) in uncertainty if $\frac{\partial V}{\partial \alpha}$ is concave (convex) in the random variable z, that is, if $\frac{\partial^3 V}{\partial z^2 \partial \alpha}$ is negative (positive). Assuming

[7] This definition of increasing uncertainty should be differentiated from the popular conception of increasing risk, which generally involves an increasing mean and may or may not involve an increase in the spread around the mean.

differentiability, it can be shown that this condition corresponds to the following:[8]

$$\frac{\partial^3 V}{\partial z^2 \partial \alpha} = \frac{\partial^3 c}{\partial \alpha \partial \mu^2}\left(\frac{\partial \mu^*}{\partial z}\right)^2 + \frac{\partial^2 c}{\partial \alpha \partial \mu}\frac{\partial^2 \mu^*}{\partial z^2} \leq (\geq)\, 0 \tag{7}$$

We interpret Equation 7 under the assumption that μ^* is increasing and concave in z—it is optimal to abate more, but decreasingly so, the larger the damage parameter, or $\frac{\partial \mu^*}{\partial z} \geq 0$ and $\frac{\partial^2 \mu^*}{\partial z^2} \leq 0$. Since abatement is ultimately bounded at 1, μ^* must be concave in the long run[9]. Under these assumptions, the response of R&D to increasing climate uncertainty depends crucially on the sign of two terms in Equation 7: (1) the impact of technological change on the marginal costs of abatement, $\frac{\partial^2 c}{\partial \alpha \partial \mu}$, and (2) the impact of technological change on the convexity of the abatement cost curve, $\frac{\partial^3 c}{\partial \alpha \partial \mu^2}$.

4.2 Results

Several insights emerge from this theoretical exercise. First, there is no absolute answer as to whether R&D should increase or decrease in uncertainty; the response of optimal R&D to uncertainty depends critically on the specifics of the R&D program. This follows from the fact that neither of the key terms in Equation 7 can be signed unambiguously. For example, while the standard assumption is that R&D will reduce the marginal cost of abatement, Program 3 increases marginal costs for some levels of abatement. Furthermore, Program 2 decreases the convexity of marginal abatement, while Program 3 increases the convexity of marginal abatement.

Adding nuance to this point, it is not possible for *any* R&D program that reduces the cost of abatement, regardless of its parameterization, to unambiguously increase with all forms of increasing uncertainty. For every R&D program, there exists some form of increasing uncertainty under which optimal R&D decreases. This notion is formalized in the following proposition. The proof is presented

[8] A derivation is provided in the appendix.
[9] A sufficient condition for concavity of μ is for damages and costs to be quadratic in μ, and damages linear in z.

in the appendix.[10]

Proposition 1 *Consider the planning problem given in (5). Assume that R&D (weakly) reduces the cost of abatement ($c(\mu, \alpha_H) - c(\mu, \alpha_L) \leq 0 \; \forall \alpha_H > \alpha_L$) and that full abatement can be achieved ($c(1, \cdot) < \infty$). Then the optimal amount of R&D will decrease with some increases in uncertainty, regardless of the program.*

One intuitive interpretation of this proposition follows from the fact that abatement is bounded above by full abatement. If damages are severe enough, full abatement (or something close to full abatement) will be optimal. For even more severe damages, any increase in realized damages will have no (or negligible) impact on optimal abatement. If damages are in this range, then increased R&D will not provide an environmental benefit in the sense of allowing for higher cost-effective abatement. It will only decrease the costs of what is essentially a fixed level of abatement. For mean-preserving spreads that increase the severity of damages in the tail of the distribution, but lower chance of getting into the tail, there will only be this second benefit and a lower probability of needing it. This reduces the expected benefits from R&D. We will return to this point in the stylized application in Sections 5 and 6.

Note that the converse of Proposition 1 does not hold. There are R&D programs that decrease unambiguously in uncertainty. For example, if costs and damages are quadratic in abatement, then optimal R&D into Program 1 unambiguously decreases in increasing uncertainty.[11]

Thus far, we have discussed the ambiguity of the results. We now ask, what attributes of a program will make it more likely to be a hedge against uncertainty? Note from Equation 7 that if R&D decreases the convexity[12] of the abatement cost function $\left(\frac{\partial}{\partial \alpha}\left[\frac{\partial^2 c}{\partial \mu^2}\right] < 0\right)$, then $\frac{\partial V}{\partial \alpha}$ may be concave for some z,

[10] Note that we are assuming that abatement costs are separable from damages. If we instead used a framework similar to that in Gollier et al. (2000), it is possible that the proposition would not hold for high levels of prudence.

[11] Let $c = (\mu - \alpha)^2$ and $D = z(1-\mu)^2$. Then $\mu^* = \frac{\alpha + z}{1+z}$,$V = (1-\alpha)^2 \frac{z}{1+z}$,$V_{\alpha zz} = 4\frac{(1-\alpha)}{(1+z)^3} \geq 0$ for $z \geq 0$

[12] In a loose sense. It is well known that the second derivative of a function is not a good measure of convexity. Thus, it is possible that R&D could increase convexity (measured as $\frac{c''}{c'}$) but decrease the second derivative, and vice-versa.

implying that R&D may increase in uncertainty and therefore be considered a hedge against increasing uncertainty. Considering the three programs discussed above, Program 1 is fairly neutral with respect to convexity, Program 2 decreases convexity, and Program 3 increases convexity. In a loose sense, then, Equation 7 argues that Program 2 is the most likely of the three to serve as a hedge against increasing uncertainty. Taking the point further, Equation 7 argues that, again in a loose sense, the more an R&D program is weighted toward technologies associated with high levels of abatement and away from those associated with low levels of abatement, the greater the chance that optimal R&D will increase in uncertainty because such programs make the abatement cost curve less convex.

5 A stylized Application Using DICE

In this section and the next, we attempt to gain more practical insight into the forces and dynamics that shape the types of R&D programs that may be a hedge against uncertainty and the types of uncertainty that may be hedged against with R&D. This section introduces a modified version of the well-known DICE model in which we incorporate R&D and climate uncertainty.[13] Section 6 discusses and interprets the results of the model.

To be clear, the application here is largely stylistic, meant to explore and understand. It is not meant to derive defensible numerical estimates of optimal R&D, but rather to augment more detailed, but deterministic studies. Along these lines, the approach includes a number of simplifications. R&D only occurs in the near-term (there is not a continuous R&D path) and climate uncertainty resolves at precisely the same time that the benefits from R&D accrue. And, at a more fundamental level, DICE itself is a very compact model of an enormously complex system.

[13] See Nordhaus and Boyer (2000) for the full description of the model.

5.1 The DICE Model

DICE is a global optimal growth model extended to include interactions between economic activities and the climate. The model covers the period from 1995 to beyond 2150 in ten-year periods. In each period, output is divided between consumption and investment in new capital, consistent with the standard optimal growth framework. DICE adds to this framework by allowing for emissions of GHGs into the atmosphere as part of the production process. The accumulation of these GHGs affects welfare by increasing global temperatures, which, in turn, reduce production. In order to mitigate this effect, an abatement level can be chosen each period, which reduces emissions through substitution below what would otherwise occur for a given production level. While abatement has obvious benefits, it is costly, reducing the amount of output available for consumption or investment in every period. The objective of the model is to maximize the discounted sum of utility over time, where utility is based on consumption. The optimal abatement path reflects a balance between benefits and costs.

Below are two equations from the model that are pertinent to the changes that will be made to incorporate R&D and uncertainty. (We do not present the full DICE model here. Readers unfamiliar with the DICE framework are encouraged to consult Nordhaus and Boyer (2000) for a more thorough treatment.)

$$Q_t = \frac{1}{1+\theta_1 T_t + \theta_2 T_t^2}(1 - b_1\mu_t^{b_2})A_t K_t^{\gamma} L_t^{(1-\gamma)} \tag{8}$$

$$E_t = (1-\mu_t)\sigma_t A_t K_t^{\gamma} L_t^{(1-\gamma)} \tag{9}$$

Equation 8 shows the output relationship in the model. Output each period, Q_t, is based on inputs of labor, L, and capital, K, modified by a technological change parameter A. Abatement, μ, reduces available output through the term $(1 - b_1\mu^{b_2})$, where $b_1 = .03$ and $b_2 = 2.15$ are parameters of the model.[14] Hence, the cost of abatement, consistent with the treatment in the previous sections, is $c(\mu) = b_1\mu^{b_2}$. The accumulation of GHGs impacts production through the temperature effects in the term $\frac{1}{1+\theta_1 T+\theta_2 T^2}$, where θ_1 and θ_2 are parameters of the model. Equation 9 shows the emissions relationship of the model. Absent abatement, emissions each period, E_t are based linearly on unadjusted production, $AK_t^{\gamma}L_t^{(1-\gamma)}$, through the emissions-output ratio parameter σ. This parameter changes over time exogenously. Abatement reduces emissions by a fraction μ.

5.2 Incorporating R&D

The R&D decision is made at the beginning of the first period, in 1995, before all other actions take place. The costs of R&D are given by $g(\alpha)$, where α represents the level of technological advance, as in above sections. These costs are subtracted from the initial capital stock as follows

$$K_1 = K - g(\alpha) \tag{10}$$

The benefits of R&D, technological advance, accrue starting in 2045. Hence, $c_t(\mu) = c_t(\alpha, \mu)$ for $t \geq 2045$, and remain unchanged for $t < 2045$. (Recall that DICE includes exogenously specified general technological change affecting the parameter A_t in Equation 8. This remains in the model for all t.) As we discuss below, for simplicity, uncertainty in climate damages is also resolved after 50 years. The time lag between effort and results is intended to capture the idea that it takes time to transition from the results of R&D to the point of widespread technology deployment.

The lack of an opportunity for additional R&D after uncertainty is resolved is an important conceptual simplification of the model. It represents the idea that by the time uncertainty is resolved, it will take too long for additional R&D to bear fruit if quick action is needed. This simplification should

[14] As noted in Section 3, production available for consumption and investment is reduced since a portion of output must be expended on abatement and therefore unavailable for consumption or investment.

be kept in mind when interpreting the model results. In particular, if a second decision were available, it would most likely put downward pressure on all of the types of near-term R&D.

The costs of R&D are given by $g(\alpha) = \frac{\kappa\alpha^2}{1-\alpha}$, where κ is a constant calibrated to each program individually. Along with simplicity, this functional form has two desirable qualities. First, it exhibits decreasing returns to scale in R&D. Second, it ensures that R&D will not bring about zero-cost full abatement.

We explore each of the three R&D programs from Section 3 individually.[15] For Program 1, the abatement cost function is modified to $c(\mu) = b_1(\mu - \alpha)^{b_2}$. For Program 2, the abatement cost function is modified to $c(\mu) = (1-\alpha)b_1\mu^{b_2}$. For Program 3, emissions from Equation 9 are multiplied by $(1-\alpha)$, thus reducing the emissions-to-output ratio, σ. The R&D cost constant, κ, was calibrated for each of the three R&D programs individually. In all cases, the calibration was conducted so that, using the standard deterministic parameters of DICE, optimal R&D would be roughly the same across the three programs.[16]

The cost coefficient for Program 1 and Program 3 was arbitrarily calibrated so that the cost of a 10 percent shift ($\alpha = 0.1$) is one percent of GDP (see Baudry, 2000). This assumption leads to $\kappa_1 = \kappa_3 = 18$. To maintain consistent R&D responses in the deterministic setting, $\kappa_2 = 1$.[17] Sensitivity analysis on the cost coefficient leads us to believe that the qualitative insights we obtain regarding the directional response of R&D to climate uncertainty are largely unaffected by changed cost coefficients.

[15] Note that the results are very similar when considering an optimal portfolio of the three programs (see the working paper, Baker, Clarke, and Weyant, 2003).

[16] To be clear, this is not the same as putting the three programs on "equal" footing. Because the three programs are functionally different, R&D productivity and optimal R&D outside of the calibration point may be markedly different. This is an inherent complication in calibrating functionally different innovation functions.

[17] Note that an argument could be made that κ_2 is reasonably lower than κ_3 in order to capture the greater opportunities for continued improvement in the less mature technologies targeted under Program 2. It is often conjectured that there is a "learning curve" associated with new technologies that tends to be very steep at the beginning and flatten out with time and experience. This implies that making improvements in mature technologies may be a great deal more expensive than making improvements in radical new technologies that have received considerably less attention.

5.3 Incorporating Uncertainty

The impact of a changed climate on human welfare is a fundamental uncertainty for climate change planning, whether for R&D or emissions policy. The less deleterious the implications of increased atmospheric GHG concentrations, the lower the appropriate near-term response. But this uncertainty has not been resolved to date, and, in fact, remains a matter of much debate.

In the DICE model, the impact on human welfare is captured as a translation of temperature increases into reduced production (see Equation 8 above), with the parameters θ_1 and θ_2 determining the degree of negative impact.

Uncertainty in climate damages is captured here by considering θ_2 to be a random variable with two possible outcomes, θ^L and θ^H, representing low damages and high damages. We take a stochastic programming approach, setting θ_2 equal to the standard value in DICE, $\theta_2 = 0.0035$, until 2045. At this point, learning takes place, and the world splits off into two paths. θ_2 may take on either the low-damage or the high-damage values, θ^L and θ^H, where the probability distribution maintains the standard value from DICE as the mean. The model is solved using GAMS/MINOS.

The use of discrete distributions to represent uncertainty presents fundamental conceptual problems. The method chosen to approximate a continuous distribution with a discrete distribution may have serious ramifications for results.[18] Of particular importance, different discrete distributions derived from the same underlying continuous distribution may have fundamentally different uncertainty characteristics—one distribution may be a mean-preserving-spread of the other, therefore exhibiting more uncertainty. Additionally, some discretization methods will involve extreme values, while others will be more centered. As will be seen in section 6, the inclusion of extreme values can have significant implications.

We use two-point distributions to underline these issues, and model increasing uncertainty using two

[18] See Smith (1993) for a discussion of the performance of four discretization methods.

distinct cases: (1) an "increasing probability" case and (2) an "increasing damages" case. Tables 1 and 2 below give the key values for each case. For the "increasing probability" case, θ^H is held constant, but its probability is increased. The value of low damages θ^L is adjusted to maintain a constant mean. Four scenarios are considered: a deterministic, and three uncertain scenarios. In the first of these uncertain scenarios, θ^H is set to .042, which implies a great depression sized loss of GDP for a 2.5° C increase in temperature. The probability of high damages for this scenario, $p = .018$, is calibrated from the estimates from Nordhaus and Boyer (2000), which in turn are based on an expert survey from Nordhaus (1994).[19] The third uncertain scenario assigns a probability of 8.3 percent to θ^H. To put this in perspective, the average probability assigned by the scientists in the survey by Nordhaus (as opposed to the economists) was 12 percent.

For the "increasing damage" case, θ^L is held constant, and θ^H is increased to capture increasing uncertainty. The probability of high damages is adjusted to maintain a constant mean. Six scenarios are considered in this case, labeled by the value of θ^H. The scenarios were chosen to illustrate the non-monotonic response of optimal R&D to increasing uncertainty. In particular, we include the scenario with $\theta^H = .057$, because this is the point at which full abatement becomes optimal even in the absence of R&D, allowing us to explore the intuition from Proposition 1. The highest damage scenario here, with $\theta^H = .1$, implies about a 38 percent loss of GDP for a 2.5° C increase in temperature.

Probability of high damage	–	1.8%	5.0%	8.3%
Value of high damage	0.0035	0.042	0.042	0.042
Value of low damage	0.0035	0.0028	0.0015	0

Table 1: Parameters for the Increasing Probability scenarios

[19] A group of experts were asked to assess the probability of a great-depression sized catastrophe (damages equal to about 20 percent of GDP) given a 2.5 C degree rise in temperature. Nordhaus combined the survey responses, then updated it with new information available since 1994, to get a probability of 1.8 percent. In our framework this implies that $\theta_2 = .042$ with probability 1.8 percent.

Probability of high damage	–	7.7%	2.1%	1.3%	0.91%	0.73%
Value of high damage	0.0035	.012	0.036	.057	0.08	0.1
Value of low damage	0.0035	0.0028	0.0028	0.0028	0.0028	0.0028

Table 2: Parameters for the Increasing Damage scenarios

6 Results of the Stylized Application

Now we turn to the response of optimal R&D to increasing climate uncertainty in the stylized empirical model. Figure 5 shows the results for the "increasing probability" case. The left-hand panel shows the results in terms of the parameter α. The right-hand panel shows the results in billions of U.S. dollars. Because of arbitrary calibration of the R&D model, the absolute amounts are less important than the relative response of each program to different levels of uncertainty. And, in this, the figure bears out the indications from the theoretical analysis: R&D in Program 1 is fairly flat (the cost curve is roughly quadratic, thus R&D has very little impact on the convexity of the cost curve), R&D in Program 2 is increasing in uncertainty, and R&D in Program 3 is decreasing in uncertainty.[20] This dynamic supports the notion that R&D into very-low emissions technologies can serve as a hedge against climate uncertainty. Investment into emissions reductions on currently economic technologies, while it may be an optimal strategy in the absence of uncertainty, is less and less attractive as uncertainty increases. This is because, at higher damage levels, this kind of R&D will have a lower payoff, as it does not focus on reducing the costs of large-scale abatement.

The "increasing damage" case tells a different story, as illustrated in figure 6. While R&D into Programs 1 and 2 initially increases in uncertainty, all three R&D programs decrease in uncertainty when larger damages ($\theta^H > .057$) are considered. This result may appear counter-intuitive, but it is driven by the logic of Proposition 1. As noted earlier, when $\theta^H = .057$, full abatement is optimal even

[20]Sensitivity analysis indicates that a change in the cost-of-R&D coefficient, κ, does not have a qualitative impact on these results. The percentage change in R&D stays about the same for different values of κ. Note, however, that the optimal amount of R&D is highly sensitive to this coefficient. Interestingly, the optimal amount to be spent on R&D is fairly insensitive.

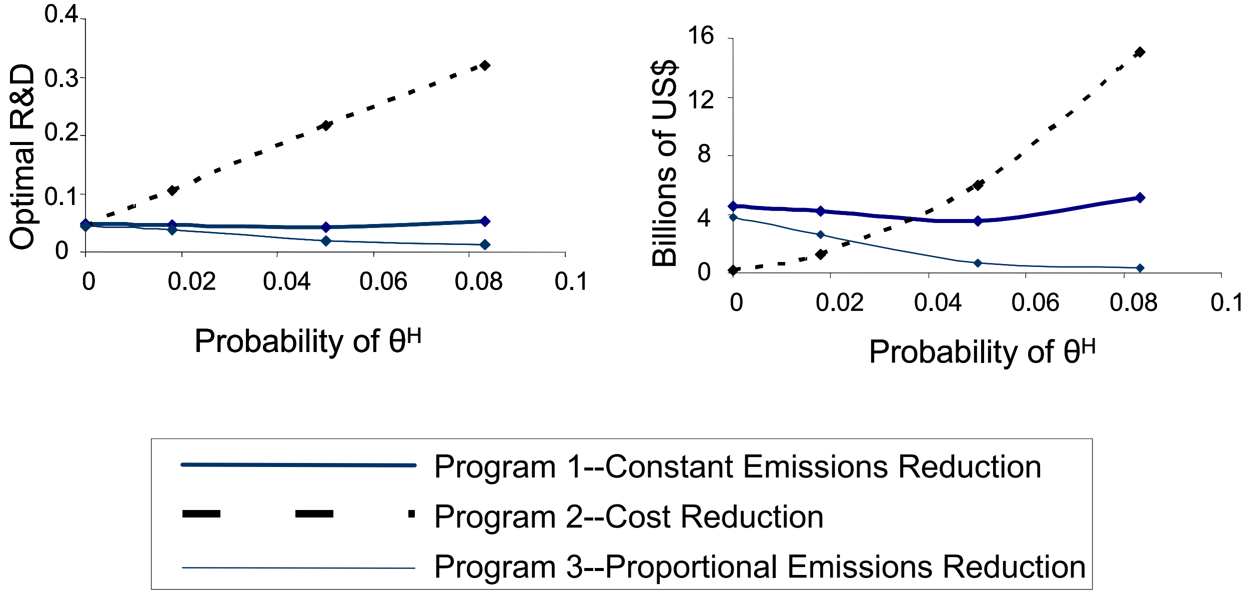

Figure 5: Optimal R&D as uncertainty increases under the assumption of increasing probability. The left-hand panel is in terms of percentage shift, α. The right-hand panel is in billions of US dollars invested.

in the absence of technical change. As θ^H increases above .057, the probability of θ^H decreases in order to preserve the mean of the distribution. Thus, the probability of full abatement (i.e. the probability of realizing the maximum benefit from R&D) decreases as θ^H increases. Thus, the expected benefit of R&D is decreasing in uncertainty, even for Program 2, the program which focuses most directly on the technologies associated with full abatement. Since full abatement is optimal even in the absence of technical change, then the presence of R&D has no impact on the optimal emissions level, it only reduces costs. In this case, there is a cost-side benefit (with decreasing probability), but no environmental-side benefit.

The argument for investing in R&D as a hedge is normally interpreted as one of whether we should be targeting technologies that would be valuable for a quick and dramatic reduction in emissions, perhaps to zero emissions. This argument appears to be valid if the concern is a higher probability of bad outcomes that call for the very-high abatement levels. However, if the concern regards larger and larger catastrophes, all calling for full abatement, with lower and lower probabilities of occurrence, then the optimal response to increasing uncertainty actually calls for lower R&D into these technologies.

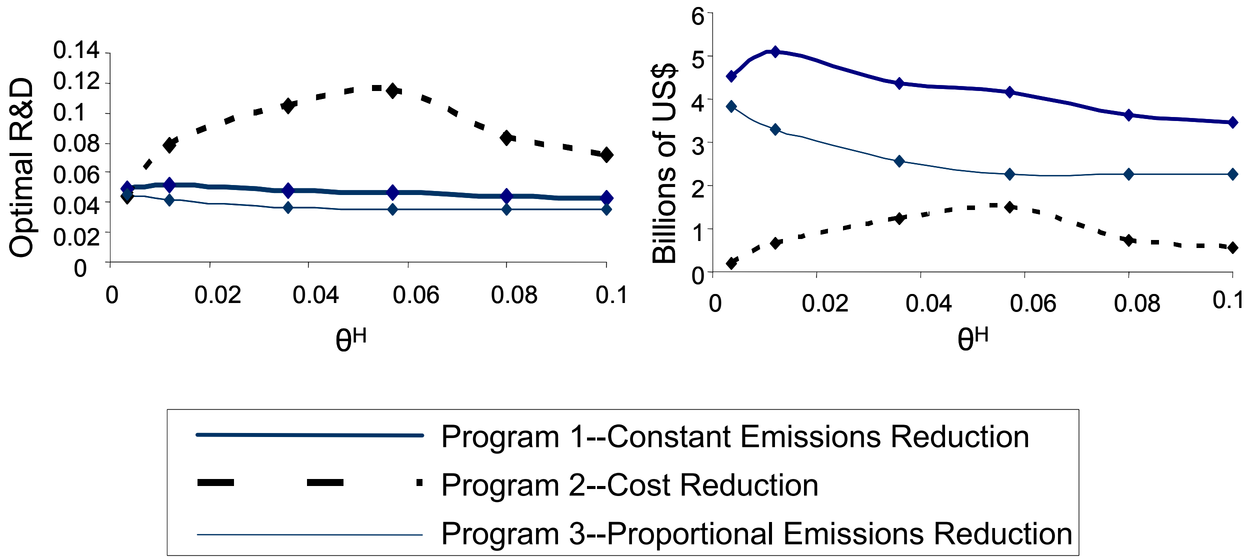

Figure 6: Optimal R&D as uncertainty increases under the assumption of increasing damages. The left-hand panel is in terms of percentage shift, α. The right-hand panel is in billions of dollars invested.

Stepping back, the results point toward a potentially productive alternative view of the relationship between uncertainty and optimal R&D. The response of optimal R&D to changes in our perception of future uncertainties—increasing uncertainty, decreasing uncertainty, or even mean-changing uncertainty—may be best explored in the light of shifting probability masses rather than mean-preserving spreads. At a general and simple level, the results indicate that whether optimal R&D into a technology will increase with a change in uncertainty depends on whether the chance of using the technology increases or decreases. This perspective is attractive because it allows consideration of increasing risk outside of the very strict economic interpretation of a mean-preserving spread. In common parlance, an increase in risk is often associated with a changed spread *and* a changed mean, limiting the practical application of conclusions based on the strict economic interpretation.

7 Conclusion

This paper has explored the response of optimal climate change technology R&D to increasing climate uncertainty. The central question is whether optimal R&D increases with increasing uncertainty, which would imply that it is a hedge against uncertainty.

The primary theme of this paper is that there is no absolute answer to the main question it addresses. Whether optimal R&D increases or decreases with increasing uncertainty depends on the character of existing technology, the manner in which R&D alters technology, and the nature of increasing uncertainty. The theoretical analysis pointed to this result, and it was borne out by the stylized, empirical analysis in which R&D increased or decreased in uncertainty depending on the R&D program and the nature of the increasing uncertainty. Taking the point further, the theoretical analysis showed that, under the assumption that full abatement is possible, no R&D program serves as a hedge against all forms of uncertainty. At the same time, however, the results indicated that R&D focused on very-low emissions technologies can serve as a hedge against increasing uncertainty, particularly in response to catastrophes that are not large enough to provoke full abatement in the absence of technology. Finally, the paper suggests that at a fundamental level, the issue is perhaps less one of increasing uncertainty in the strict paradigm of a mean-preserving spread, and more one of whether changed uncertainty increases the probability of needing a technology on which we might conduct R&D.

Several caveats are in order based on the limitations of the analysis. First, increasing uncertainty was expressed in the strict economic sense of a mean-preserving spread. This is an important assumption, as common discourse on climate "risk" is often devoid of precise definitions of "risk", and often assumes an increasing mean, not just an increasing spread. Second, the analysis assumes that no R&D is conducted after climate uncertainty is resolved. This restricts the meaningful domain of the results and insights to situations where emissions reduction actions in response to the resolution of uncertainty are required on a timescale that is shorter than the timescale of technology development and deployment. In general, the opportunity to do R&D after learning takes place will reduce the amount of R&D done in the short term. Third, the stylized application includes a number of simplifications, including relatively arbitrary paramaterizations of R&D, the assumption that all climate uncertainty is resolved in 2045, that R&D can only be conducted right now and not continually over the century, and that, at the broadest level, the DICE model with its standard parameterization is a meaningful conception of reality. These assumptions do not impact the qualitative results but have serious ramifications for any quantitative interpretation of the results. Fourth, the empirical work only considers R&D that alters abatement costs but does not impact technology more generally, although this is not a difficult generalization. For example, an R&D program which shifts the production function down by an absolute amount would be independent from both the expected value and the risk of damages. Finally, it is an open question where current R&D efforts lie—there is no reason to believe that they are in any way an optimal starting point from which we might consider deviations. Future work might address any or all of these limitations.

A Appendix

The social planning problem is solved by working backwards, beginning with optimal second-period abatement based on the level of technical advance and the resolution of climate uncertainty. Optimal second-period abatement, μ^*, is defined by the first order condition

$$\frac{\partial c}{\partial \mu} + \frac{\partial D}{\partial \mu} = 0 \tag{11}$$

Using this, the social planning problem can now be written as in (6), where $V(\alpha, z) = c(\mu^*; \alpha) + D(\mu^*; z)$ is an indirect welfare function, representing the minimum impact on second period welfare (abatement costs plus damages), given α and z. Optimal abatement, given by μ^*, is a function of both α and z. The first order condition for the social planning problem is

$$g'(\alpha) = -\,\mathrm{E}_z\left[\frac{\partial V}{\partial \alpha}\right] \tag{12}$$

Any change in the distribution of z that increases the right-hand-side will increase the optimal level of α. A concave function decreases in uncertainty.[21] Thus, $\mathrm{E}_z\left[\frac{\partial V}{\partial \alpha}\right]$ will decrease in uncertainty if $\frac{\partial V}{\partial \alpha}$ is concave in z , thus the right-hand side will increase in uncertainty if $\frac{\partial V}{\partial \alpha}$ is concave in z. Assuming differentiability, $\frac{\partial V}{\partial \alpha}$ is concave in z if $\frac{\partial^3 V}{\partial z^2 \partial \alpha} \leq 0$. To calculate this quantity , we first take the derivative of V with respect to α.

$$\frac{\partial V}{\partial \alpha} = \frac{\partial c}{\partial \alpha} \tag{13}$$

by the envelope theorem. Now take derivatives with respect to z:

$$\frac{\partial^2 V}{\partial \alpha \partial z} = \frac{\partial^2 c}{\partial \alpha \partial \mu}\frac{\partial \mu^*}{\partial z} \tag{14}$$

$$\frac{\partial^3 V}{\partial \alpha \partial z^3} = \frac{\partial^3 c}{\partial \alpha \partial \mu^2}\left(\frac{\partial \mu^*}{\partial z}\right)^2 + \frac{\partial^2 c}{\partial \alpha \partial \mu}\frac{\partial^2 \mu^*}{\partial z^2} \tag{15}$$

A.1 Proof of Proposition 1

Proof. In order to prove the proposition without assuming differentiability, we use a more general version of the Rothschild-Stiglitz Theorem (see Corollary 1 in Athey (2000) or Baker (2002)), that says that α unambiguously increases in risk if and only if $V(\alpha_H, z) - V(\alpha_L, z)$ is concave in z for all $\alpha_H \geq \alpha_L$. Thus, we need to show that $V(\alpha_H, z) - V(\alpha_L, z)$ is not everywhere concave in z. Specifically, we show that $V(\alpha_H, z) - V(\alpha_L, z)$ must be convex as z gets large.

1. Without lack of generality we can normalize z so that damages are 0 at $z = 0$. This implies that $V(\alpha_H, 0) - V(\alpha_L, 0) = 0$.

[21] Jensen's inequalitiy is a special case of this general result. See Rothschild & Stiglitz (1970).

2. By assumption, $c(\mu, \alpha_H) - c(\mu, \alpha_L) \leq 0$ for $\alpha_H \geq \alpha_L$. Optimality implies that

$$V(\alpha_H, z) = c\left(\mu^H, \alpha_H\right) + D\left(\mu^H, z\right) \leq c(\mu, \alpha_H) + D(\mu, z) \;\; \forall \mu \tag{16}$$

where μ^H represents optimal abatement when $\alpha = \alpha_H$. In particular

$$V(\alpha_H, z) \leq c\left(\mu^L, \alpha_H\right) + D\left(\mu^L, z\right) \tag{17}$$

$$= c\left(\mu^L, \alpha_H\right) - c\left(\mu^L, \alpha_L\right) + c\left(\mu^L, \alpha_L\right) + D\left(\mu^L, z\right) \tag{18}$$

$$\leq c\left(\mu^L, \alpha_L\right) + D\left(\mu^L, z\right) = V(\alpha_L, z) \;\; \forall \alpha_H > \alpha_L \tag{19}$$

Thus

$$V(\alpha_H, z) - V(\alpha_L, z) \leq 0 \; \forall z, \alpha_H > \alpha_L \tag{20}$$

3. Combining 1 and 2 implies that $V(\alpha_H, z) - V(\alpha_L, z)$ is (weakly) decreasing in z at $z = 0$.

4. We show that $V(\alpha_H, z) - V(\alpha_L, z)$ is bounded from below, and therefore not everywhere concave.

$$V(\alpha_H, z) - V(\alpha_L, z) = c\left(\mu^H, \alpha_H\right) + D\left(\mu^H, z\right) - \left[c\left(\mu^L, \alpha_L\right) + D\left(\mu^L, z\right)\right] \tag{21}$$

$$\geq c\left(\mu^H, \alpha_H\right) + D\left(\mu^H, z\right) - \left[c\left(\mu^H, \alpha_L\right) + D\left(\mu^H, z\right)\right] \tag{22}$$

$$= c\left(\mu^H, \alpha_H\right) - c\left(\mu^H, \alpha_L\right) \geq -c(1, 0) > -\infty \tag{23}$$

Inequality (22) follows from optimality. The conclusion in (23) follows since c is bounded by assumption.

■

Acknowledgement 1 *Excellent research assistance was provided by Kwame Adu-Bonnah. Baker ac-*

knowledges partial funding from the Global Climate and Energy Project, Stanford University, and from a Faculty Research Grant, University of Massachusetts, Amherst.

References

Athey, S. (2000). Characterizing properties of stochastic objective function. Manuscript, Stanford University, Palo Alto.

Baker, E. (2002). *Increasing Informativenss and Increasing Risk: Equivalence Theorems*. PhD thesis, Stanford University, Palo Alto.

Baker, E. (2003). Uncertainty and learning in a strategic environment: Global climate change. Submitted to Resource and Energy Economics.

Baudry, M. (2000). Joint management of emission abatement and technological innovation for stock externalities. *Environmental and Resource Economics*, 16:161–183.

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

Clarke, L. and Weyant, J. (2002). Modeling induced technological change: An overview. In Grubler, A., Nakicenovic, N., and Nordhaus, W., editors, *Technological Change and the Environment*. Resources for the Future, Washington, D.C.

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

Gollier, C., Jullien, B., and Treich, N. (2000). Scientific progress and irreversibility: An economic interpretation of the 'precautionary pricipal'. *Journal of Public Economics*, 75:229–253.

Goulder, L. and Mathai, K. (2000). Optimal CO_2 abatement in the presence of induced technological change. *Journal of Environmental Economics and Management*, 39(1):1–38.

Goulder, L. and Schneider, S. (1999). Induced technological change and the attractiveness of CO_2

abatement policies. *Resource and Energy Economics*, 21:211–253.

Jaffe, A., Newell, R., and Stavins, R. (2001). Technological change and the environment. Discussion Paper 00-47REV, Resources for the Future, Washington, D.C.

Keller, K., Bolker, B., and Bradford, D. (2004). Uncertain climate thresholds and optimal economic growth. Forthcoming in Journal of Environmental Economics and Management.

Kolstad, C. (1996). Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. *Journal of Environmental Economics and Management*, 31:1–18.

Manne, A. and Richels, R. (2002). The impact of learning-by-doing on the timing and cost of CO_2 abatement. Technical Report Working Paper, AEI-Brookings Joint Center for Regulatory Studies.

Manne, A. S. (1996). Hedging strategies for global carbon dioxide abatement: A summary of poll results - EMF 14 subgroup analysis for decisions under uncertainty. Energy Modeling Forum, Stanford University.

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

Nordhaus, W. D. (1994). Expert opinion on climatic change. *American Scientist*, 82:45–51.

Nordhaus, W. D. (2002). Modeling induced innovation in climate change policy. In Grubler, A., Nakicenovic, N., and Nordhaus, W. D., editors, *Technological Change and the Environment*, pages 182–209. RFF and IIASA, Washington D.C. and Laxenburg, Austria.

Nordhaus, W. D. and Boyer, J. (2000). *Warming the World: Economic Models of Global Warming*. MIT Press, Cambridge, MA.

Pizer, W. (1999). The optimal choice of climate change policy in the presence of uncertainty. *Resource* *and Energy Economics*, 21(3/4):255–287.

Popp, D. (2002). ENTICE: Endognenous technological change in the DICE model of global warming. Syracuse University, Submitted to *Journal of Environmental Economics and Management*.

Rothschild, M. and Stiglitz, J. (1970). Increasing risk I: A definition. *Journal of Economic Theory*, 2:225–243.

Rothschild, M. and Stiglitz, J. (1971). Increasing risk II: Its economic consequences. *Journal of Economic Theory*, 3:66–84.

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

Smith, J. E. (1993). Moment methods for decision analysis. *Management Science*, 39:340–358.

Sue Wing, I. (2003). Induced technical change and the cost of climate policy. Technical Report 102, MIT Joint Program on the Science and Policy of Global Change.

Ulph, A. and Ulph, D. (1997). Global warming, irreversibility and learning. *The Economic Journal*, 107:636–650.

van der Zwaan, B., Gerlagh, R., Klaassen, G., and Schrattenholzer, L. (2002). Endogenous technological change in climate change modeling. *Energy Economics*, 24:1–19.

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