

**Sunk Costs and Sunk Benefits in Environmental Policy:
I. Basic Theory**

by

Robert S. Pindyck

MIT-CEEPR 95-003WP

March 1995

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

SEP 05 1996

Library

SUNK COSTS AND SUNK BENEFITS IN ENVIRONMENTAL POLICY: I. BASIC THEORY*

by

Robert S. Pindyck
Massachusetts Institute of Technology
Cambridge, MA 02139

This draft: March 15, 1995

Abstract: The standard framework in which economists evaluate environmental policies is cost-benefit analysis, so that policy debates usually focus on expected costs and benefits, or on the choice of discount rate. But this framework can be misleading when there is uncertainty over future outcomes, when there are irreversibilities, and when policy adoption can be delayed. This paper shows how “economic” uncertainty (i.e., uncertainty over future costs and benefits of reduced environmental degradation) and “ecological” uncertainty (i.e., uncertainty over the evolution of an ecosystem) interact with two kinds of irreversibilities — sunk costs associated with an environmental regulation, and “sunk benefits” of avoided environmental degradation — to affect optimal policy timing and design.

JEL Classification Numbers: Q28, L51, H23

Keywords: Environmental policy, irreversibilities, cost-benefit analysis, uncertainty, opportunity cost, option value, global warming.

*This research was supported the National Science Foundation under Grant No. SES90-22823, by M.I.T.'s Center for Energy and Environmental Policy Research, by the M.I.T. Program on Global Climate Change, and by the U.S. Department of Energy under Agreement No. DE-FG41-92R110447. My thanks to Avinash Dixit, Charles Kolstad, Ian Parry, and Martin Weitzman for helpful comments and suggestions, and to Eileen Brooks and Stephen McNamara for their superb research assistance.

1. Introduction.

The standard framework in which economists evaluate environmental policies is cost-benefit analysis. Consider, for example, a carbon tax to reduce global warming. By distorting relative prices, this policy would impose an expected flow of costs on society in excess of the government tax revenues it generates. Presumably, it also yields an expected flow of benefits to society. Households and firms would burn less fuel, less CO₂ would accumulate in the atmosphere, global mean temperatures would not rise as much, and the damage caused by higher temperatures would be correspondingly smaller. The standard framework would recommend this policy if the present value of the expected flow of benefits exceeds the present value of the expected flow of costs. Any debate among economists would likely be over the expected costs and benefits, or over the choice of discount rate.¹

This standard framework, however, ignores three important characteristics of most environmental problems and the policies designed to respond to them. First, there is uncertainty over the future costs and benefits of adopting a particular policy. With global warming, for example, there is considerable uncertainty over how much average temperatures are likely to rise with or without reduced CO₂ emissions, and also over the economic impact of higher temperatures. Second, there are usually important irreversibilities associated with environmental policy. These irreversibilities can arise with respect to environmental damage itself, but also with respect to the costs of adapting to policies to reduce the damage. Third, the adoption of an environmental policy is rarely a now or never proposition. In most cases it is feasible to delay action and wait for new information.

Environmental policy design thus has something in common with the design of an irreversible investment policy. When a firm makes an irreversible investment, it gives up the possibility of waiting for new information that might affect the desirability or timing of the expenditure; it cannot disinvest should market conditions change adversely. This lost option value is an opportunity cost that must be included as part of the cost of the investment

¹There is a large literature on the choice of discount rate and its dependence on risk and rates of taxation. For an overview, see Lind (1984). For a more recent study especially relevant to environmental policy, see Weitzman (1994).

when doing net present value (NPV) calculations. Investment in a project is warranted only when the present value of a project's expected cash flows exceeds the project's cost, at least by an amount equal to the value of keeping the investment option alive. Hence a project with a conventional NPV that is positive may in fact be uneconomical.²

In the case of environmental policy, the implications of irreversibility and uncertainty are more complicated. The reason is that there are two kinds of irreversibilities, and they work in opposite directions. First, policies aimed at reducing ecological damage impose *sunk costs* on society. These sunk costs can take the form of discrete investments; for example, coal-burning utilities might be forced to install scrubbers, or firms might have to scrap existing machines and invest in more fuel-efficient ones. Or they can take the form of flows of expenditures, e.g., a price premium paid by a utility for low-sulfur coal. In either case, such sunk costs create an opportunity cost of adopting a policy now, rather than waiting for more information about ecological impacts and their economic consequences. This opportunity cost biases traditional cost-benefit analysis *in favor* of policy adoption. As with irreversible investment decisions, the sunk costs associated with policy adoption can make it preferable to wait rather than adopt the policy now.

Second, environmental damage can be partially or totally irreversible. For example, increases in GHG concentrations are long lasting. Even if radical policies were adopted in the future to drastically reduce GHG emissions, these concentrations (which have a natural decay rate of about a half percent per year) would take many years to fall. In addition, the damage to various ecosystems from higher global temperatures (or from acidified lakes and streams, or from the clear-cutting of forests) can be permanent. This means that adopting a policy now rather than waiting has a *sunk benefit*, i.e., a negative opportunity cost. This negative opportunity cost biases traditional cost-benefit analysis *against* policy adoption. Hence it may be desirable to adopt a policy now, even though the traditional analysis declares it uneconomical.

This point regarding irreversible environmental damage was made nearly two decades

²For an introduction to and overview of the literature on irreversible investment, see Dixit (1992) and Pindyck (1991). For a more detailed treatment, see Dixit and Pindyck (1994).

ago by Arrow and Fisher (1974), Henry (1974), and Krutilla and Fisher (1975), and has been elaborated upon by Hanemann (1989), and Fisher and Hanemann (1990), among others. However, little has been done to evaluate its practical importance, or to determine whether (for particular environmental problems and policies) the sunk costs of policy adoption outweigh the sunk benefits of environmental preservation. At issue is whether these irreversibilities are important, and if so, what their overall effect is.

The answer is likely to depend on the nature and extent of uncertainty. In general, two different types of uncertainty are relevant. The first, which I call “economic uncertainty,” is uncertainty over the future costs and benefits of environmental damage and its reduction. In the context of global warming, for example, even if we knew how large a temperature increase to expect from any particular increase in GHG concentrations, we would not know what cost society will bear as a result. The reason is that we are unable to predict how a particular temperature increase would affect agricultural output, land use, etc.³ Likewise, even if we could predict the increase in acidity in lakes and rivers from NOX emissions, the impact of this increase on fish and other organisms is uncertain and hence so is the social cost. Indeed, for most environmental problems there is inherent uncertainty over the future social cost of the environmental degradation, and thus over the social benefit of any policy response.

Second, most environmental problems involve at least some degree of what I call “ecological uncertainty,” i.e., uncertainty over the evolution of various ecosystems and environmental variables. For example, while we could specify a policy target for the rate of GHG emissions over the next forty years, we would not know what the resulting levels of atmospheric GHG concentrations will be at different points in time, nor would we know what the average global equilibrium temperature increase would be, and how that increase would vary regionally. And even given assumptions about economic growth in different parts of the world, predicting GHG emissions (in the absence or presence of policy intervention) is difficult,

³The impacts on agricultural output, for example, is likely to vary substantially from region to region, and, as Mendelsohn, Nordhaus, and Shaw (1994) have pointed out, are likely to be mitigated by regional shifts in land utilization.

and subject to considerable uncertainty.⁴ Similarly, we are unable to accurately predict how particular levels of NOX emissions will affect the future acidity of lakes and rivers, or the viability of the fish that live in them.

This paper develops a series of models to help elucidate the implications of irreversibility and uncertainty for environmental policy. These models are intended to show how the sunk costs of policy adoption and the sunk benefits of environmental preservation interact with the two types of uncertainty described above. An objective is a framework that can be applied to problems like global warming and acid rain, and that will yield at least rough estimates of the bias inherent in a traditional cost-benefit analysis.

Two recent studies that are closely related to this paper deserve mention. First, Kolstad (1992) developed a three-period model to study the implications of cost-benefit uncertainty for the adoption of an emissions-reducing policy that can involve sunk costs. In his model, the accumulated stock of pollutant is permanent. Emissions can be reduced in the first or second periods, and between these periods there is a reduction in uncertainty over the net benefits from a lower stock of pollutant. He shows that if there is no sunk cost of policy adoption, then the faster the rate of learning, the lower the first-period emission level should be. This is essentially a version of the result of Arrow and Fisher (1974) and Henry (1974); because the stock of pollutant is permanent, society should pollute less now if there is uncertainty over the future damage from the pollutant. But he also shows that if emission control is irreversible (i.e., the costs of policy adoption are at least partly sunk), then the effect of uncertainty on the initial level of emissions is ambiguous.⁵

Second, Hendricks (1992) has developed a continuous-time model of global warming similar to the one in this paper. As I do, he studies the timing of policies to irreversibly

⁴For a forecasting model of CO₂ emissions with an explicit treatment of forecast uncertainty, see Schmalensee, Stoker, and Judd (1995). For general discussions of the uncertainties inherent in the analysis of global warming, see Cline (1992) and Solow (1991), and, for a good overview of these uncertainties and some of their policy implications, see Jacoby and Prinn (1994).

⁵In related work, Hammitt, Lempert, and Schlesinger (1992) use a two-period model to study implications of uncertainty for adoption of policies to reduce GHG emissions, and show that under some conditions it may be desirable to wait for additional information. Also, Kolstad (1994) examines GHG emission policy in the context of a growth model with uncertainty and learning, and finds that temporary emission-reduction policies dominate permanent ones.

reduce emissions, allowing for a (partially) irreversible accumulation of the pollutant. The particular form of uncertainty he considers is over a parameter linking the global mean temperature increase to the atmospheric GHG concentration, and he allows for learning by assuming that the uncertainty over this parameter falls over some fixed period of time. He focuses on how the speed of learning affects the timing of policy adoption.⁶

Unlike these studies, I assume that while information is continually arriving over time, there will always remain uncertainty over the future evolution of key environmental variables, and over the future costs and benefits of policy adoption. I then focus on how the sunk costs and sunk benefits of policy adoption interact with these uncertainties in affecting the timing and design of policy. I do this by developing a series of models that are increasingly general in terms of the policy choices they allow, and in terms of the types of uncertainty they incorporate.

The next section begins by laying out the basic analytical framework for all of the models, and shows how the design and timing of environmental policy in the presence of uncertainty can be treated as an optimal stopping problem. Then, in Section 3, I begin with a simplified problem in which an environmental policy can be adopted either today or at a specified future time, and adoption entails a flow of sunk costs (some portion of which may be recoverable if the policy is later removed). Limiting decisions to only two points in time makes it possible to obtain and analyze optimal policies with minimal use of mathematics, and helps to clarify the basic mechanism through which uncertainty (in this case over the future benefits from reduced pollution) and the two kinds of irreversibilities affect the timing of policy adoption.

In Section 4, it is assumed that policy adoption can occur at any time, but there is no uncertainty. Nonetheless, we will see that if the social cost of environmental damage is growing over time (e.g., due to population growth), it may be desirable to delay policy adoption and thereby discount its cost, even when a standard cost-benefit analysis would indicate immediate adoption.

⁶Although he does not do so, Hendricks could also use his model to study the implications of the degree of irreversibility of environmental damage. This could be done by varying the parameter that describes the rate of natural GHG removal from the atmosphere.

Sections 5 to 8 develop models of policy adoption under uncertainty in continuous time. In Section 5, it is assumed once again that a policy to reduce emissions can be adopted at any time, but now economic uncertainty is introduced. I first consider the case in which policy adoption implies reducing emissions to zero, and then the case in which the size of the reduction can be chosen optimally at the time of adoption. In addition, I examine the policy timing problem for both linear and convex economic benefit functions. In Section 6, I allow for *gradual* emission reductions, again in the presence of economic uncertainty. Section 7 examines the implications of ecological uncertainty. Finally, Section 8 presents a general model that includes both economic and ecological uncertainty.

2. Analytical Framework.

In a traditional cost-benefit analysis of environmental policy, the problem typically boils down to whether or not a particular policy should be adopted. When irreversibilities are involved, the more appropriate question is *when* (if ever) a particular policy should be adopted. In other words, adopting a policy today competes not only with never adopting the policy, but also with adopting the policy next year, adopting it in two years, and so on. The policy problem becomes one of optimal timing.

This is exactly analogous to the optimal irreversible investment problem. The NPV criterion tells us to invest in a project when the present value of the expected payoff stream from the project just exceeds the present value of the cost. But this criterion implicitly compares investing today with never investing. When the investment is at least partly irreversible and the firm can delay, the correct comparison is investing today versus waiting for more information to arrive and perhaps investing at some point in the future. The problem is again one of optimal timing.

This analogy also holds when policies can be reversed, either totally or partially. When an environmental policy is adopted, it leads to flows of expenditures by firms and consumers that are in part sunk, i.e., irreversible. If, after a number of years, we learn that the environmental problem that the policy seeks to remedy is actually much less of a problem than had been expected, the policy might be reversed, perhaps limiting these flows of sunk costs. With

ongoing uncertainty, deciding when to reverse the policy is also an optimal timing problem. (The reason is that the decision to reverse also has a sunk cost component, and could lead to regret if some years later new information leads to greater concern with the environmental problem.) Likewise, a firm that can exit a market and recover part of its sunk costs must decide on the timing of exit, as well as the timing of entry.⁷

In order to get at the basic issues and obtain results that are reasonably easy to interpret, it is best to begin with a model that captures the basic stock externality associated with many environmental problems in as simple a way as possible, while still allowing us to capture key sources of uncertainty. Let M_t be a state variable that summarizes one or more stocks of environmental pollutants. For example, M might be the average concentration of CO₂ in the atmosphere, the acidity level of a lake, the number of plant or animal species in a forest, or the concentrations of a mix of pollutants that make up urban smog. Let E_t be a flow variable that controls M_t . For example, E might be the rate of CO₂ or SO₂ emissions, or a rate of deforestation. We will assume that absent some policy intervention, E_t follows an exogenous (and possibly stochastic) trajectory. Ignoring uncertainty for the time being, the evolution of M_t is then given by:

$$dM/dt = \beta E(t) - \delta M(t), \quad (1)$$

where δ is the natural rate at which the stock of pollutant dissapates over time.

So far I have described a simplified version of a basic diffusion model used by Nordhaus (1991) to compare costs and benefits of policies to reduce greenhouse gas (GHG) emissions. That model supplements eqn. (1) with an adjustment process for temperature:

$$dT/dt = \alpha[\mu M(t) - T(t)], \quad (2)$$

where T is the increase in mean temperature from GHGs, M is atmospheric GHG concentration from industrial activity, and α is a delay parameter. Associated with a higher T is a (global) economic cost resulting from, among other things, land loss due to a rising sea

⁷This issue of investment timing when there is partial reversibility is explored in Dixit and Pindyck (1994, 1995).

level, and reduced agricultural output due to climate change. I am simplifying the Nordhaus model by dropping the variable T along with eqn. (2), and associating an economic cost directly with M .⁸ Also, note that eqn. (1) is deterministic; later I will introduce ecological uncertainty by generalizing this equation so that M follows a controlled diffusion process.

I will assume that the flow of social cost (i.e., negative benefit) associated with the stock variable M_t is specified by a function $B(M_t, \theta_t)$, where θ_t is a variable that shifts over time, perhaps stochastically, to reflect changes in tastes and technologies. For example, if M is the GHG concentration, shifts in θ might reflect new agricultural techniques that reduce the cost of a higher M , or alternatively, demographic changes that raise the cost. Or, if M is the concentration of an industrial pollutant collecting in a body of water and damaging the fish population, shifts in θ might reflect changes in the price of fish, or changes in the value of sport fishing.

Later we will explore the implications of different functional forms for $B(M_t, \theta_t)$, but for the time being I will assume that it is simply:

$$B(M_t, \theta_t) = -\theta_t M_t. \quad (3)$$

In the next section, uncertainty over the future costs and benefits of policy adoption will be introduced by letting θ follow a stochastic process.

We will be interested in the design and timing of policies that change the evolution of E_t , typically making it smaller than it would be otherwise. The implications of uncertainties and irreversibilities are easiest to understand by focussing on policies that are introduced at a specific point in time, and that have a long-term impact on the evolution of E_t . (In later sections of the paper we will also consider policies that are introduced gradually and can be modified from time to time.) Consider a policy introduced at time T that changes the evolution of E_t for $t \geq T$. Such a policy would presumably impose a flow of costs on society, some portion of which will be sunk. We will denote the present value (at time T) of the expected flow of *sunk* costs associated with this policy by $K(E_T, \omega)$, where ω is a vector

⁸Because of the linearity of eqns. (2) and (1), both could have been retained in all of the analyses that follow (with an economic cost associated with T , or with both M and T). The basic qualitative results would remain the same.

of policy characteristics. For example, ω might describe an absolute reduction in E_t , or a reduction in the expected rate of growth in E_t compared to what it would be absent the policy.⁹

Again in the interests of clarity, we will start out by making some very restrictive assumptions about the evolution E_t , and about possible policies to change E_t . Specifically, we will assume that until a policy is adopted, the rate of emissions E_t stays at the constant initial level E_0 . Policy adoption involves a once-and-for-all reduction in E_t to some new and permanent level E_1 , with $0 \leq E_1 \leq E_0$. I will also begin by assuming, again for simplicity, that the social cost of adopting this policy is completely sunk, and its present value at the time of adoption is a convex function of the size of the emission reduction, which I denote by $K(E_1)$. Note that the policy might entail a *flow* of sunk costs over time (e.g., expenditures for insulation on all new homes). All that matters is that adopting the policy implies a commitment to this flow of costs, so that we can replace the flow with its present value at the time of adoption.

It is easy to generalize this problem so that the policy involves gradual reductions in the emission rate, so that only part of the costs of adopting the policy are sunk, and so that the policy can be reversed in the future, leading to a partial recovery of the sunk cost K . I will discuss such generalizations later, although they do little to change the basic results. At this point, note that the policy problem involves a choice of timing, and (at the time of adoption) a choice of how much to reduce emissions.

The policy objective is to maximize the net present value function:

$$W = \mathcal{E}_0 \int_0^{\infty} B(M_t, \theta_t) e^{-rt} dt - \mathcal{E}_0 K(E_1) e^{-r\tilde{T}}, \quad (4)$$

subject to eqn. (1). Here, \tilde{T} is the (in general, unknown) time that the policy is adopted, $E_0 - E_1$ is the amount that emissions are reduced, \mathcal{E}_0 denotes the expectation at time $t = 0$,

⁹For example, we might have an emission level E_t that, absent a policy intervention, will grow stochastically according to:

$$dE_t = \alpha_E E_t dt + \sigma_E E_t dz_E.$$

Then, a policy might involve a one-time reduction in E_T (thereby reducing the expected value of E_t for all $t \geq T$), or it might involve a reduction in α_E , the expected rate of growth of E_t .

and r is the discount rate. This is an optimal stopping problem. Specifically, we must find a rule that tells us when it is optimal to commit to spending K to reduce E_t , given the (possibly stochastic) dependence of M_t on E_t , and given the stochastic evolution of θ_t . As we will see, this stopping rule will not give us the particular time at which to adopt the policy, but rather conditions that M_t and θ_t must meet for adoption to be optimal. Hence the time of adoption, \tilde{T} , is a random variable.

We will approach this optimal stopping problem in steps, first considering the case of no uncertainty, and, for simplicity, assuming that the policy involves reducing E_t to zero. Then, we will introduce uncertainty over the evolution of θ_t , over M_t , and finally over both. We will also generalize the problem so that emission reduction can be done gradually, rather than all at once, and so that the policy can be reversed and emissions allowed to rise again. First, however, it will be helpful to look at a much simpler version of this problem of policy timing under uncertainty.

3. A Simplified Problem.

The usual evaluation of an environmental policy compares the outcome of adopting the policy today with the outcome of never adopting it. As a simple first step, let us introduce a third possibility — the option of adopting the policy at some *fixed time* T in the future. In other words, one can decide to adopt the policy today, one can wait until time T and then, after evaluating the situation, decide whether or not to adopt the policy, or one can decide now to never adopt the policy. Since this last possibility is clearly suboptimal, the choice boils down to adopting the policy now or waiting until time T to decide.

We will begin by assuming that if the policy is adopted, emissions are reduced from E_0 to zero. Hence the cost of policy adoption is simply a number, K , which we will assume is complete sunk. Later in this section we will consider the possibility of reducing E to some level $E_1 > 0$. We will also examine the adoption decision when the policy is partially reversible — if adopted now, it can be reversed at time T , leading to the recovery of some fraction ϕ of the sunk cost K .

For this problem to be interesting, we need to introduce some source of uncertainty. I

will assume that there is cost-benefit uncertainty but not ecological uncertainty, i.e., there is uncertainty over the evolution of θ_t but not over the evolution of M_t . To keep matters as simple as possible, I will assume that θ_T will equal $\underline{\theta}$ or $\bar{\theta}$ with equal probability, with $\underline{\theta} < \bar{\theta}$ and $\frac{1}{2}(\underline{\theta} + \bar{\theta}) = \theta_0$, the current value of θ . I will also assume that θ does not change after time T . Finally, I will consider the following decision rule that applies *if* we wait until time T : Adopt the policy if and only if $\theta_T = \bar{\theta}$. (I will choose parameter values so that this is indeed the optimal policy at time T .) Later we will see that this very limited form of uncertainty over the evolution of θ can be generalized considerably without affecting the basic qualitative results.

By solving eqn. (1), we can determine M_t as a function of time. Suppose the policy is adopted at time T , so that $E_t = E_0$ for $t < T$ and $E_t = 0$ for $t \geq T$. Then it is easy to confirm that M_t is given by:

$$M_t = \begin{cases} (\beta E_0/\delta)(1 - e^{-\delta t}) + M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T \\ (\beta E_0/\delta)(e^{\delta T} - 1)e^{-\delta t} + M_0 e^{-\delta t} & \text{for } t > T \end{cases} \quad (5)$$

where M_0 is the initial value of M_t . If the policy is never adopted, the first line of eqn. (5) applies for all t , so that M_t asymptotically approaches $\beta E_0/\delta$. If the policy is adopted at time 0, then $M_t = M_0 e^{-\delta t}$.

First, suppose that the policy is never adopted. Then, denoting the value function in this case by W_N :

$$\begin{aligned} W_N &= - \int_0^{\infty} \theta_0 M_t e^{-rt} dt \\ &= -\theta_0 \int_0^{\infty} [(\beta E_0/\delta)(1 - e^{-\delta t}) + M_0 e^{-\delta t}] e^{-rt} dt \\ &= -\frac{\theta_0 M_0}{(r + \delta)} - \frac{\beta E_0 \theta_0}{r(r + \delta)} \end{aligned} \quad (6)$$

Next, suppose the policy is adopted at time $t = 0$. Then $E_t = 0$ always, and the value function is:

$$W_0 = -\frac{\theta_0 M_0}{r + \delta} - K. \quad (7)$$

Note that a conventional cost-benefit analysis would recommend adoption of the policy if the net present value $W_0 - W_N$ is positive.

Let us introduce some numbers so that we can compare these two alternatives. I will assume that the present value of the cost to society of policy adoption is \$2 billion. I will set $r = .04$, $\delta = .02$, $\beta = 1$ (i.e., emissions are completely absorbed into the ecosystem), $E_0 = 300,000$ tons/year, and $\theta_0 = \$20/\text{ton}/\text{year}$.¹⁰ In what follows, I will also assume that $\underline{\theta} = \$10/\text{ton}/\text{year}$, and $\bar{\theta} = \$30/\text{ton}/\text{year}$.

Given these numbers, $\beta E_0 \theta_0 / r(r + \delta) = \2.5 billion. Since the conventionally measured NPV of policy adoption is $W_0 - W_N = \beta E_0 \theta_0 / r(r + \delta) - K = \0.5 billion, it would appear desirable to adopt the policy now.

Suppose that instead we wait until time $t = T$ and then adopt the policy only if $\theta_T = \bar{\theta}$. Denoting the value function in this case by W_T , using eqn. (5), and noting that the probability that $\theta_T = \bar{\theta}$ is .5, we have:

$$W_T = -\frac{\theta_0}{r + \delta} \left(M_0 + \frac{\beta E_0}{r} \right) + \frac{\beta E_0}{r(r + \delta)} (\theta_0 - \frac{1}{2}\underline{\theta}) e^{-rT} - \frac{1}{2} K e^{-rT}. \quad (8)$$

Is it better to adopt the policy at time $t = 0$ or wait until T ? We can decide by comparing W_0 to W_T :

$$\begin{aligned} \Delta W_T &= W_T - W_0 \\ &= K(1 - \frac{1}{2}e^{-rT}) - \frac{\beta E_0 \theta_0}{r(r + \delta)} (1 - e^{-rT}) - \frac{\beta E_0 \underline{\theta}}{2r(r + \delta)} e^{-rT}. \end{aligned} \quad (9)$$

The first term on the right-hand side of eqn. (9) is the present value of the net expected cost savings from delay; the sunk cost K is initially avoided, and there is only a .5 probability that it will have to be incurred at the later time T . Hence this term represents the opportunity cost of adopting the policy now rather than waiting. The second and third terms on the RHS of eqn. (9) are the present value of the expected increase in social cost from environmental damage due to delay. The second term is the the social cost of additional pollution between now and time T that results from delaying policy adoption. The last term in the equation — the probability that $\theta_T = \underline{\theta}$, times the present value of the social cost of

¹⁰I am implicitly assuming that the discount rate r is the real risk-free rate of interest, so a value of .04 is reasonable. A value of .02 for δ is high for the rate of natural removal of atmospheric GHGs (a consensus estimate would be closer to .005), but is low for acid concentrations in lakes.

additional pollution over time when $\theta_T = \underline{\theta}$ and $E_t = E_0$ for $t \geq T$ — is the expected pollution cost from time T onwards. Thus the last two terms together represent an opportunity “benefit” of adopting the policy now rather than waiting.

We can therefore rewrite eqn. (9) as:

$$\Delta W_T = F_C - F_B,$$

where

$$F_C = K(1 - \frac{1}{2}e^{-rT}) \quad (10)$$

is the opportunity cost of adopting the policy now rather than waiting, and

$$F_B = \frac{\beta E_0 \theta_0}{r(r + \delta)}(1 - e^{-rT}) + \frac{\beta E_0 \underline{\theta}}{2r(r + \delta)}e^{-rT} \quad (11)$$

is the opportunity “benefit” of adopting now rather than waiting. Note that the larger is the decay rate δ , i.e., the more reversible is environmental damage, the smaller is this benefit, and hence the greater is the incentive to delay. An increase in the discount rate, r , increases the cost F_C and reduces the benefit F_B , and thus also increases the incentive to delay.

In general, we can decide whether it is better to wait or adopt the policy now by calculating F_C and F_B . For our numerical example, we will assume (arbitrarily) that the fixed time Y is 10 years. Substituting this and the other base case parameter values into eqns. (10) and (11) gives $F_C = \$1.330$ billion and $F_B = 0.824 + 0.419 = \$1.243$ billion. Hence $\Delta W_T = F_C - F_B = \0.087 , so it is better to wait rather than adopt the policy now. In this case the sunk cost of current adoption slightly outweighs the sunk benefit.

We assumed in these calculations that if we delayed the adoption decision until time T , it would be optimal to adopt the policy if $\theta_T = \bar{\theta}$, but not if $\theta_T = \underline{\theta}$. To check that this is indeed the case, we can calculate the smallest value of θ_T for which policy adoption at time T is optimal. Since there is no possibility of delay after time T , this is just the value of θ for which $W_0 - W_N$ is zero. Denoting this value by $\hat{\theta}_T$ and using eqns. (7) and (6), we see that it is given by:

$$\hat{\theta}_T = r(r + \delta)K/\beta E_0. \quad (12)$$

For our base case parameter values, $\hat{\theta}_T = \$16/\text{ton}/\text{year}$. Hence it would indeed be optimal to adopt the policy at time T if $\theta_T = \bar{\theta} = 30$, but not if $\theta_T = \underline{\theta} = 10$.

Also, we assumed that policy adoption meant reducing E to zero. We could have instead considered policies to reduce E by some fixed percentage, or considered what the optimal amount of reduction should be. However, we have assumed that the social cost function $B(M_t, \theta_t)$ is linear in M_t , and because M_t depends linearly on E (see eqn. (5)), the benefit of a marginal reduction in E is independent of the level of E . Suppose, in addition, that the cost of reducing E is proportional to the size of the reduction. Then if it is optimal to reduce E at all, it will be optimal to reduce it to zero, so that the optimal timing of the reduction is independent of the size of the reduction. This will not be the case, however, if the social cost function is convex in M_t , and/or the cost of emission reduction is a convex function of the size of the reduction. We will discuss this further below.

Irreversibility, Uncertainty, and a “Good News Principle.”

We have assumed that the cost of adopting the policy is completely sunk, but the benefit (in terms of reduced environmental damage) is only partially sunk (because $\delta > 0$). Continuing with our numerical example, we can get further insight into the effects of irreversibility and uncertainty by varying the degree to which the policy benefit is sunk, and by varying the amount of uncertainty over θ_T .

First, suppose that the rate at which pollutants are naturally removed from the ecosystem is slower than assumed earlier — specifically, that the parameter δ is .01 instead of .02. Note that F_C will equal \$1.330 billion as before, but now $F_B = 0.989 + 0.503 = \$1.492$ billion, so that $\Delta W_T = -\$0.162$ billion. In this case the greater irreversibility of environmental damage makes the sunk benefit of current adoption greater than the sunk cost, so that it is better to adopt the policy now.

Second, let us increase the variance of θ_T (while keeping its expectation the same) by setting $\underline{\theta}$ and $\bar{\theta}$ equal to 0 and 40 respectively, instead of 10 and 30. This change has no effect on the sunk cost of adopting now, because there is still a .5 probability that at time T we will regret having made the decision to spend K and adopt the policy; note that F_C

is \$1.330 billion as before. However, this increase in variance reduces the sunk benefit of immediate adoption by reducing the social cost of additional pollution for $t > T$ under the “good” outcome (i.e, the outcome that $\theta_T = \underline{\theta}$). Setting δ equal to its base case value of .02, we now have $F_B = .824 + 0 = \$0.824$ billion, so that $\Delta W_T = 1.330 - 0.824 = \0.506 billion, which is much larger than before. Even if we increase the irreversibility of environmental damage by setting δ equal to .01, $F_B = .989$, so that $\Delta W_T = \$0.341$ billion and it is still optimal to wait.

This important result is an example of Bernanke’s (1983) “bad news principle,” although here we might call it a “good news principle.” It is only the consequences of the outcome $\theta_T = \underline{\theta}$ — an outcome that is good news for society, but bad news for the *ex post* return on policy-induced installed capital — that drive the net value of waiting. The consequences of the “bad” outcome, i.e., the outcome $\theta_T = \bar{\theta}$, make no difference whatsoever in this calculation.

This good news principle might seem counterintuitive at first. Given the long-lasting impact of environmental damage, one might think that the consequences of the high social cost outcome (i.e., the outcome $\theta_T = \bar{\theta}$) should affect the decision to wait and continue polluting instead of adopting the policy now. But since the *expected value* of θ_T remains the same as we increase the variance, the value of waiting depends only on the regret that is avoided under the good (low social cost) outcome. Increasing the variance of θ_T increases the regret that society would experience under the good outcome, and thereby increases the incentive to wait. Shortly we will see that this result holds even if we allow for a policy adopted at time 0 to be reversed at time T , with partial recovery of the sunk cost K .

Generalizing the Process for θ_t .

Our assumptions about the evolution of θ_t are extremely restrictive — it can have only two possible values at time T , and does not change after time T . The same basic results will hold, however, if θ_t follows a more general process. For example, suppose θ_t follows the geometric Brownian motion:

$$d\theta = \alpha\theta dt + \sigma\theta dz. \tag{13}$$

Once again, we will assume that the policy can be adopted now, or at some fixed time T in the future. Should we wait until time T , we will adopt the policy if and only if $\theta_T > \hat{\theta}_T$, where $\hat{\theta}_T$ defines the optimal stopping rule. Hence we first need to find the smallest $\hat{\theta}_T$ such that we would want to adopt the policy if θ_T were above this value. Since there is no option to delay the adoption decision beyond time T , this is straightforward. As before, we use eqns. (7) and (6) to find $W_0 - W_N$, and then set this equal to zero, with $\hat{\theta}_T$ replacing θ_0 . Since $\mathcal{E}_0\theta_t = \theta_0 e^{\alpha t}$, we get

$$\hat{\theta}_T = (r - \alpha)(r - \alpha + \delta)K/\beta E_0. \quad (14)$$

Note that if $\alpha = 0$, this is identical to eqn. (12) above, and for our base case parameters, $\hat{\theta}_T$ is again equal to \$16.

To determine whether it is optimal to wait rather than adopt the policy now, we again calculate $\Delta W_T = F_C - F_B$. Let p_T equal the probability that $\theta_T < \hat{\theta}_T$. Since θ_T is lognormally distributed and we know the value of $\hat{\theta}_T$, we can find p_T numerically (or from a table of the lognormal distribution). Next, let $\underline{\theta} \equiv \mathcal{E}_0(\theta_T | \theta_T < \hat{\theta}_T)$. (This can likewise be calculated by a simple numerical integration.) Then, eqns. (10) and (11) for F_C and F_B again apply, except that the $\frac{1}{2}$ in each is replaced by p_T .

It is easy to see that a change in δ , or a change in the variance of θ_T (in this case a change in σ), will have the same general effects as before. Consider an decrease in δ . As before, F_C will not change, but F_B will increase since the stock of pollutant is longer lasting. If the increase in F_B is sufficiently large, it may shift the optimal timing from $t = 0$ to $t = T$. Next, consider what happens if σ increases. Again, this has no effect on F_C . But an increase in the variance of θ_T implies a decrease in $\underline{\theta}$, so that F_B falls. This is the “good news principle” again — an increase in uncertainty over θ_T tends to delay policy adoption by exacerbating the “downside” risk associated with current policy-related sunk costs.

Allowing for Policy Reversal.

So far we have assumed that a policy to reduce emissions to zero could be adopted at time 0 or time T , but once adopted, the policy would remain in place indefinitely. Now we will see how the timing decision changes when a policy that has been adopted at time 0 can

be reversed at time T . We will assume that upon reversal, a fraction ϕ of the policy-induced cost K can be recovered. A partial recovery of K would be possible, for example, if K was at least in part the present value of a flow of sunk costs that could be terminated upon reversal of the policy. (Of course, the investment decisions of firms and consumers in response to a policy adopted at time 0 would be altered by the awareness that there was some probability of policy reversal at time T . For example, consumers and firms would probably delay some of their emission-reducing investments until they learned, at time T , whether the policy was going to be reversed. But this is consistent with the theory; it simply makes the fraction ϕ larger than it would be without such awareness.)

We will assume as we did earlier that θ_T can have one of two values, $\underline{\theta}$ or $\bar{\theta}$. We will also assume that the parameter values are such that if the policy was not adopted at $t = 0$, it would be adopted at $t = T$ if and only if $\theta_T = \bar{\theta}$. However, if the policy is adopted at $t = 0$, would we want to remove it at time T if $\theta_T = \underline{\theta}$? Clearly, this will depend on ϕ .

As before, we will let W_0 denote the value function when we adopt the policy at time 0. But now, W_0 must include the value of society's option (a *put* option) to reverse the policy at time T and recover ϕK . We will again let W_T denote the value function when we wait and only adopt the policy if $\theta_T = \bar{\theta}$. (In this simple two-period framework, we do not allow for the possibility of reversing the policy *after* time T .) To determine W_0 , we need to know the trajectory for M_t when the policy is adopted at $t = 0$ and reversed at $t = T$. Solving eqn. (1), that trajectory is given by:

$$M_t = \begin{cases} M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T \\ (\beta E_0 / \delta) [1 - e^{-\delta(t-T)}] + M_0 e^{-\delta t} & \text{for } t > T \end{cases} \quad (15)$$

Now we can determine the minimum value of ϕ for which it would be economical to reverse the policy at $t = T$ should $\theta_T = \underline{\theta}$. It is economical to remove the policy as long as the present value of the cost of continued emissions is less than the recoverable cost ϕK , i.e., as long as:

$$(\beta E_0 \underline{\theta} / \delta) \int_T^\infty [1 - e^{-\delta(t-T)}] e^{-r(t-T)} dt < \phi K, \quad (16)$$

which implies

$$\phi_{\min} > \frac{\beta E_0 \underline{\theta}}{r(r + \delta)K}. \quad (17)$$

For our numerical example, with $\underline{\theta} = \$10/\text{ton}/\text{year}$, $\phi_{\min} = .838$. Thus if $\phi < .838$, the option to reverse the policy at time T has no value, and our earlier results still hold.

Suppose $\phi = .90$, so that the policy would indeed be reversed if $\theta_T = \underline{\theta}$. Although W_T is still given by eqn. (8), using eqn. (15), W_0 is now given by:

$$W_0 = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_0 \underline{\theta}}{2r(r + \delta)} e^{-rT} + \frac{1}{2} \phi K e^{-rT} - K. \quad (18)$$

The second and third terms on the right-hand side of this equation represent the value of option to reverse the policy at time T .

Using eqns. (8) and (18), we find that $\Delta W_T = W_T - W_0$ is now given by:

$$\Delta W_T = \frac{\beta E_0 \theta_0}{r(r + \delta)} (1 - e^{-rT}) + K \left[1 - \frac{1}{2}(1 + \phi)e^{-rT} \right]. \quad (19)$$

The first term on the right-hand side of (19) is the ‘‘opportunity benefit’’ of early policy adoption, which we have denoted by F_B , and the second term is the opportunity cost of early adoption, denoted by F_C . Comparing this to eqn. (11), we see that F_B no longer has the term in $\underline{\theta}$, because now if $\theta_T = \underline{\theta}$ the policy will be reversed.

Let us return to our numerical example, with $\phi = .9$. In this case,

$$\Delta W_T = -\$824.2 \text{ million} + \$726.4 \text{ million} = -\$97.6 \text{ million},$$

so it is better to adopt the policy at time 0, rather than wait. The reason is that now F_C , the opportunity cost of adopting now, is much smaller because of the option to reverse the policy later.

Suppose we increase the variance of θ_T keeping its mean the same. As we did earlier, we will let $\underline{\theta}$ and $\bar{\theta}$ be 0 and 40 respectively, rather than 10 and 30. If $\phi = .9$, $\Delta W_T = -\$97.6$ million as before, so the policy should still be adopted now. But note that increasing the variance of θ_T reduces the minimum value of ϕ at which it is optimal to reverse the policy if $\theta_T = \underline{\theta}$. From eqn. (17), we see that now $\phi_{\min} = 0$. But this does not mean that as long as $\underline{\theta} = 0$, the policy should be adopted now for any positive value of ϕ . For example, if

$\phi = .1$, $\Delta W_T = \$438.4$ million, so it is clearly better to wait. By setting $\Delta W_T = 0$ (again with $\underline{\theta} = 0$), we can find the smallest value of ϕ for which early adoption is optimal. Using eqn. (19), that value is $\phi = .754$. For values of ϕ above .754, the put option is sufficiently valuable so that early adoption is economical.

Although $\underline{\theta}$ does not appear in eqn. (19), it is still only $\underline{\theta}$, and not $\bar{\theta}$, that affects the timing decision. The reason is that only $\underline{\theta}$ affects θ_{\min} , and hence the decision to exercise the put option should this low value of θ_T be realized. This is another example of the “good news principle” that arose earlier.

Partial Reduction in Emissions.

Before moving to a more general model in which the time of adoption can be chosen freely, we can exploit this simple framework still further by allowing for a partial reduction in emissions. As discussed earlier, this extension is of interest only if the cost of policy adoption is a convex function of the amount of emission reduction (or, alternatively, the benefit function $B(M_t, \theta_t)$ is convex in M_t). Suppose that the cost of reducing E from E_0 to $E_1 \geq 0$ is given by:

$$K = k_1(E_0 - E_1) + k_2(E_0 - E_1)^2, \quad (20)$$

with $k_1, k_2 \geq 0$. Then the cost of a 1-unit (permanent) reduction in E , given that currently $E = E_1$, is:

$$k(E) = -\frac{dK}{dE} = k_1 + 2k_2(E_0 - E_1). \quad (21)$$

The problem now is to decide when to adopt a policy, and then, at the time of adoption, to decide by how much to reduce emissions. As before, we will assume that θ_T will equal $\underline{\theta}$ or $\bar{\theta}$ with equal probability, and that θ does not change after time T . For simplicity, we will assume that once a policy has been adopted it cannot be reversed.

Previously we solved eqn. (1) to determine M_t when $E_t = E_0$ for $t < T$ and $E_t = 0$ for $t \geq T$. Now, policy adoption at time T implies $E_t = E_1 \geq 0$ for $t \geq T$. In this case, the

solution for M_t is given by:¹¹

$$M_t = \begin{cases} (\beta E_0/\delta)(1 - e^{-\delta t}) + M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T \\ (\beta E_0/\delta)(e^{\delta T} - 1)e^{-\delta t} + (\beta E_1/\delta)[1 - e^{-\delta(t-T)}] + M_0 e^{-\delta t} & \text{for } t > T \end{cases} \quad (22)$$

First, suppose we reduce E from E_0 to an arbitrary level E_1 at $t = 0$. Then the value function is:

$$W_0(E_1) = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_1 \theta_0}{r(r + \delta)} - K(E_1). \quad (23)$$

If we never adopt the policy, the value function is $W_N = -\theta_0 M_0/(r + \delta) - \beta E_0 \theta_0/r(r + \delta)$, as before. Hence the conventionally measured NPV of policy adoption is:

$$W_0(E_1) - W_N = \frac{\beta(E_0 - E_1)\theta_0}{r(r + \delta)} - K(E_1). \quad (24)$$

If we indeed adopt the policy at $t = 0$, we will choose E_1 to maximize this NPV. Using eqn. (20) for $K(E_1)$, the optimal E_1 is:

$$E_1^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\theta_0}{2k_2 r(r + \delta)}, \quad (25)$$

for $\beta\theta_0/r(r + \delta) > k_1$, and 0 otherwise. Assuming $\beta\theta_0/r(r + \delta) > k_1$ and $E_1 = E_1^*$, the NPV becomes:

$$W_0(E_1^*) - W_N = \frac{1}{4k_2} \left[\frac{\beta\theta_0}{r(r + \delta)} - k_1 \right]^2. \quad (26)$$

Note that because E_1 is chosen optimally, this NPV is never negative.

A numerical example is again helpful. We will use the same parameter values as before, i.e., $r = .04$, $\delta = .02$, $\beta = 1$, $\theta_0 = \$20/\text{ton}/\text{year}$, $\bar{\theta} = \$30/\text{ton}/\text{year}$, $\underline{\theta} = \$10/\text{ton}/\text{year}$, and $E_0 = 300,000$ tons/year. Suppose that $k_1 = 4000$ and $k_2 = .02$. (Then reducing E from 300,000 tons/year to zero would cost \$3.0 billion.) In this case, $\Delta E^* = E_0 - E_1^* = 108,333$ tons/year, $K(\Delta E^*) = \$0.668$ billion, and the NPV of immediate policy adoption is $W_0(E_1^*) - W_N = \$0.234$ billion.

Now suppose that we instead wait until time T ($= 10$ years) to decide how much (if at all) to reduce emissions. If $\theta_T = \bar{\theta}$ we will reduce emissions to \bar{E} , but if $\theta_T = \underline{\theta}$ we will

¹¹Note that M_t must now satisfy the boundary conditions $M_T = (\beta E_0/\delta)(1 - e^{-\delta T}) + M_0 e^{-\delta T}$ and $M_\infty = \beta E_1/\delta$.

reduce emissions by a smaller amount, to \underline{E} , with $\underline{E} > \bar{E}$. Using eqn. (22) for M_t and for now letting \underline{E} and \bar{E} be arbitrary, we can determine that the value function $W_T(\underline{E}, \bar{E})$ is:

$$W_T(\underline{E}, \bar{E}) = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_0 \theta_0}{r(r + \delta)}(1 - e^{-rT}) - \frac{\beta e^{-rT}}{2r(r + \delta)}(\underline{E}\theta + \bar{E}\bar{\theta}) - \frac{1}{2}K(\underline{E})e^{-rT} - \frac{1}{2}K(\bar{E})e^{-rT}. \quad (27)$$

The values of \underline{E} and \bar{E} must be chosen optimally to maximize $W_T(\underline{E}, \bar{E})$. Those values are:

$$\underline{E}^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\theta}{2k_2r(r + \delta)}, \quad (28)$$

$$\bar{E}^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\bar{\theta}}{2k_2r(r + \delta)}. \quad (29)$$

Should we reduce emissions now or wait until time T so that we can observe θ_T ? As before, we compare W_0 to W_T , but now accounting for the fact that the amount of emission reduction is determined optimally at the time of adoption, i.e., at $t = 0$ or $t = T$. To determine whether it is better to wait, we must calculate $\Delta W_T = W_T(\underline{E}^*, \bar{E}^*) - W_0(E_1^*)$. Substituting \underline{E}^* and \bar{E}^* into eqn. (27) and E_1^* into eqn. (23) gives:

$$\Delta W_T = \frac{k_1}{2k_2} \left[\frac{\beta\theta_0}{r(r + \delta)} - \frac{k_1}{2} \right] (1 - e^{-rT}) - \frac{\beta^2\theta_0^2}{4k_2r^2(r + \delta)^2} + \frac{\beta^2(\theta^2 + \bar{\theta}^2)}{8k_2r^2(r + \delta)^2} e^{-rT}. \quad (30)$$

Using eqns. (25), (28), and (29), we can calculate that for our numerical example, $E_1^* = 191,667$ tons/year, $\underline{E}^* = 295,833$, and $\bar{E}^* = 87,500$. Hence we find that $\Delta W_T = \$0.068$ billion. In this case the opportunity cost of reducing emissions immediately outweighs the sunk benefit. Therefore it is better to wait until time T , and then reduce emissions by a large amount if $\theta_T = \bar{\theta}$, but reduce them only slightly if $\theta_T = \underline{\theta}$. But note that this outcome is dependent on our choice of parameters for the cost function K . For example, if we reduce k_1 from 4000 to 1000 (so that the cost of eliminating the first ton of emissions is only \$1,000), ΔW_T becomes $-\$0.076$ billion, so that immediate policy adoption is preferred. The reason is that now a greater reduction in E is optimal for any θ (now $E_1^* = 116,667$, $\underline{E}^* = 220,833$, and $\bar{E}^* = 12,500$), so that the sunk benefit of immediately reducing E is larger, and the sunk cost smaller.

As with the simpler versions of this two-period model, the timing decision also depends on the variance of θ_T . To see this, let us increase the variance by setting $\bar{\theta}$ and $\underline{\theta}$ to 40 and 0 respectively. Now, using eqns. (25), (28), and (29) again, we see that $E_1^* = 191,667$ tons/year as before, but $\underline{E}^* = 400,000$ tons/year, $\bar{E}^* = 0$, and $\Delta W_T = \$0.504$ billion.¹² Hence the value of waiting increases. The reason is that the spread between \underline{E}^* and \bar{E}^* is now larger, so that information arriving at time T has a bigger impact on policy actions, and on the outcomes of those actions.

This simplified problem illustrates how the optimal timing of policy adoption can be affected in opposing ways by the interaction of uncertainty with each of two kinds of irreversibilities. For example, by reducing the rate at which pollutants are naturally removed from the ecosystem (making environmental damage more irreversible), we increased the sunk benefit of early policy adoption to the point where it outweighed the sunk cost. To explore this tradeoff further, and determine how it depends on different sources of uncertainty, we need to move to a more general formulation in which the time of adoption is a free choice variable.

4. Optimal Timing in a World of Certainty.

At this point it might appear that in the absence of uncertainty, irreversibilities can be disregarded, so that standard cost-benefit analysis will apply. As we will see, if the social cost of the stock of pollutant (as measured by θ in this model) is growing over time, it may be optimal to wait rather than adopting a policy now, and thereby allowing the postponement (and thus discounting) of the sunk cost K . We can explore this, and get further insight into the optimal timing of policy, by generalizing our model so that θ_t evolves continuously (but deterministically) over time, and allowing the policy to be adopted at *any* point in time, rather than either immediately or at a fixed time T as in the preceding section. Once again, for simplicity we will begin by assuming that policy adoption implies that E_t is reduced from E_0 to zero.

¹²Using eqn. (29), $\bar{E}^* = -16,667$. But we assume that negative values of E are not possible, so that E will be reduced to 0 if $\theta_T = \bar{\theta}$.

We will let θ grow at the constant rate α , so that $\theta_t = \theta_0 e^{\alpha t}$. Letting T be the time the policy is adopted (which is yet to be determined), and using eqn. (5) for M_t , the value function is given by:

$$\begin{aligned} W &= \max_T \left\{ -\frac{\beta E_0 \theta_0}{\delta} \int_0^T (1 - e^{-\delta t}) e^{(\alpha-r)t} dt - \frac{\beta E_0 \theta_0}{\delta} (e^{\delta T} - 1) \int_T^\infty e^{(\alpha-r-\delta)t} dt - K e^{-rT} \right\} \\ &= \max_T \left\{ -\frac{\beta E_0 \theta_0 [1 - e^{-(r-\alpha)T}]}{(r-\alpha)(r+\delta-\alpha)} - K e^{-rT} \right\} \end{aligned} \quad (31)$$

Maximizing with respect to T yields:

$$\frac{\beta E_0 \theta_0 e^{\alpha T}}{r + \delta - \alpha} = rK. \quad (32)$$

Since $\theta_0 e^{\alpha T} = \theta_T$, condition (32) implies that the policy should be adopted once θ reaches the following critical value:

$$\theta^* = r(r + \delta - \alpha)K / \beta E_0. \quad (33)$$

If $\theta_0 \geq \theta^*$, the policy should be adopted immediately. If $\theta_0 < \theta^*$ and $\alpha > 0$, the policy is adopted at a later time T^* , given by:

$$T^* = \frac{1}{\alpha} \log \left[\frac{r(r + \delta - \alpha)K}{\beta E_0 \theta_0} \right]. \quad (34)$$

To interpret these results, let W_0 be the value function when the policy is adopted at time 0, let W_N be the value function when the policy is never adopted, and let W_T be the value function when the policy is adopted at an *arbitrary* time T . These value functions are simply:

$$W_0 = -K, \quad (35)$$

$$W_N = -\frac{\beta E_0 \theta_0}{(r-\alpha)(r+\delta-\alpha)}, \quad (36)$$

$$W_T = -\frac{\beta E_0 \theta_0 [1 - e^{-(r-\alpha)T}]}{(r-\alpha)(r+\delta-\alpha)} - K e^{-rT}. \quad (37)$$

There are four possibilities. (1) $W_0 - W_N < 0$ and $W_T < W_N$ for all T . In this case the policy is never adopted. (2) $W_0 - W_N > 0$ and $W_T < W_0$ for all $T > 0$. In this case the

policy is adopted at $t = 0$. (3) $W_0 - W_N < 0$ but $W_T > W_N$ for some range of T . In this case the conventional NPV is initially negative, but will become positive in the future. the policy is adopted at the time T^* given by eqn. (34), which maximizes W_T . (4) $W_0 - W_N > 0$ but $W_T > W_0$ for some range of T . Now the conventional NPV is positive, but it is better to wait. Again, the policy is adopted at the T^* that maximizes W_T .


These four cases are illustrated in Figure 1. It is Case (4), however, that is most interesting, because a simple calculation that compared adopting the policy now with never adopting it would find a positive NPV. But it is better to wait; the NPV (calculated at time $t = 0$) is larger if the policy is adopted at $T^* > 0$. It may seem counterintuitive that growth in the social cost of environmental damage should result in a delay in the adoption of a policy to stop the damage. The reason for delay in this case is that waiting reduces the cost of the policy by a factor of e^{-rT} , but because θ is growing at the rate α , waiting reduces the benefit by the smaller factor of $e^{-(r-\alpha)T}$.

This result is analogous to one for the optimal timing of a capital investment under certainty. If the value of a project is non-stochastic but is growing over time, it may be best to delay investing in the project even if the NPV for immediate investment is positive. Again, this is because delaying the investment reduces the present value of its cost by more than it reduces present value of its the payoff.¹³


A numerical example is again useful. Following up on our example from the previous section, we will again assume that $r = .04$, $\delta = .02$, $\beta = 1$, $E_0 = 300,000$ tons/year, and $\theta_0 = \$20/\text{ton}/\text{year}$. We will also assume that $\alpha = .005$, and we will consider different values for K . Figure 2 shows W_0 , W_N , and W_T as functions of T for $K = \$3.5$ billion. This corresponds to Case (3); $W_0 < W_N$, but W_T eventually exceeds W_N , reaching a maximum value of - \$3.05 billion at $T^* = 50$ years. Case (4) is illustrated in Figure 3, which shows W_0 , W_N , and W_T for $K = \$2.9$ billion. Now $W_0 > W_N$, but it is better to wait. $W_0 = -\$2.90$ billion, but W_T attains a maximum value of $-\$2.86$ billion at $T^* = 12$ years.

¹³See Dixit and Pindyck (1994), pages 138–139, for a discussion. This point was first noted by Marglin (1963, Chapter 2).


(1) $W_0 < W_N, W_T < W_N$ for all T



(2) $W_0 > W_N, W_T > W_N$ for all T



(3) $W_0 < W_N, W_T > W_N$ at $T = T^*$



(4) $W_0 > W_N, W_T > W_N$ at $T = T^*$




Figure 1: Policy Timing Under Certainty.




Figure 2: Numerical Example — $\alpha = .005$, $K = \$3.5$ billion.




Figure 3: Numerical Example — $\alpha = .005$, $K = \$2.9$ billion.

5. Cost-Benefit Uncertainty.

We now generalize the model to allow for continuous stochastic fluctuations in the social benefit function, and policy adoption at any point in time. Initially, we will continue to assume that policy adoption implies reducing E to zero. This is indeed optimal if the cost of emission reduction is a linear function of the size of the reduction, and if $B(M_t, \theta_t)$ is linear in M_t . Later we will make the cost of reducing emissions a convex function of the size of the reduction, and then allow for incremental reductions in E over time.

We introduce uncertainty by letting θ follow the geometric Brownian motion:

$$d\theta = \alpha\theta dt + \sigma\theta dz. \quad (38)$$

For the time being we will assume that the cost of reducing E from its initial value, E_0 , to zero is $K = kE_0$. The problem is to find a rule for policy adoption that maximizes the net present value function of eqn. (4) subject to eqn. (38) for θ and eqn. (1) for M .

We can solve this problem using dynamic programming by defining the net present value function for each of two regions. Let $W^N(\theta, M)$ denote the value function for the “no-adopt” region (so that $E_t = E_0$), and let $W^A(\theta, M)$ denote the value function for the “adopt” region (in which $E_t = 0$). Then $W^N(\theta, M)$ must satisfy the following Bellman equation:¹⁴

$$rW^N = -\theta M + (\beta E_0 - \delta M)W_M^N + \alpha\theta W_\theta^N + \frac{1}{2}\sigma^2\theta^2 W_{\theta\theta}^N. \quad (39)$$

Likewise, $W^A(\theta, M)$ must satisfy the Bellman equation:

$$rW^A = -\theta M - \delta M W_M^A + \alpha\theta W_\theta^A + \frac{1}{2}\sigma^2\theta^2 W_{\theta\theta}^A. \quad (40)$$

These equations must be solved subject to the following boundary conditions:

$$W^N(0, M) = 0, \quad (41)$$

$$W^N(\theta^*, M) = W^A(\theta^*, M) - K, \quad (42)$$

$$W_\theta^N(\theta^*, M) = W_\theta^A(\theta^*, M). \quad (43)$$

¹⁴For an introduction to or review of the mathematical tools used in this paper, see Chapters 3 and 4 of Dixit and Pindyck (1994).

Here θ^* is the critical value of θ at or above which the policy should be adopted. Condition (41) reflects the fact that if θ is ever zero, it will remain at zero. Condition (42) is the value matching condition; it simply says that when $\theta = \theta^*$ and society exercises its option to adopt the policy, it incurs a sunk cost $K = kE_0$ and hence receives the net payoff $W^A(\theta^*, M) - K$. Condition (43) is the “smooth pasting condition;” if adoption at θ^* is indeed optimal, the derivative of the value function must be continuous at θ^* .

These differential equations and associated boundary conditions have the following solution:

$$W^N(\theta, M) = A\theta^\gamma - \frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_0 \theta}{(r - \alpha)(r + \delta - \alpha)}, \quad (44)$$

and

$$W^A(\theta, M) = -\frac{\theta M}{r + \delta - \alpha}, \quad (45)$$

where A is a positive constant to be determined, and, from boundary condition (41), γ is the positive root of the quadratic equation $\frac{1}{2}\sigma^2\gamma(\gamma - 1) + \alpha\gamma - r = 0$, i.e.,

$$\gamma = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left(\frac{\alpha}{\sigma^2} - \frac{1}{2}\right)^2 + \frac{2r}{\sigma^2}} > 1. \quad (46)$$

W^N , the present value function that applies before adoption of the policy, has three components. The first term on the right-hand side of (44) is the value of the option to adopt the policy at some time in the future. The second term is the present value of the flow of social cost resulting from the current stock of pollutant, M . (The current stock, M , decays at the rate δ , while θ has an expected rate of growth α , so the present value is $-\theta M/(r + \delta - \alpha)$.) The third term is the present value of the flow of social cost that would result if emissions continued at the rate E_0 forever. (The present value of the flow of cost from emissions E_0 now is $\beta E_0 \theta/(r + \delta - \alpha)$, but the present value of cost for emissions E_0 now and in all future periods is $\beta E_0 \theta/(r + \delta - \alpha)(r - \alpha)$.) Note that this last component of social cost is reduced by the value of the option to reduce emissions, i.e., the first term. Once the policy has been adopted, $E = 0$ and the value function W^A applies. Then the only social cost is from the current stock of pollutant.

There are still two unknowns, the constant A and critical value θ^* at which the policy

should be adopted, and they are determined from boundary conditions (42) and (43):

$$A = \left(\frac{\gamma - 1}{K} \right)^{\gamma - 1} \left[\frac{\beta E_0}{(r - \alpha)(r + \delta - \alpha)\gamma} \right]^{\gamma}, \quad (47)$$

$$\theta^* = \left(\frac{K}{\gamma - 1} \right) \left[\frac{(r - \alpha)(r + \delta - \alpha)\gamma}{\beta E_0} \right]. \quad (48)$$

We can now see how the optimal timing of policy adoption depends on the degree of uncertainty over future costs and benefits, and on other parameters. First, note that an increase in σ implies a decrease in γ and hence an increase in θ^* . As we would expect, the greater the uncertainty over the future social cost of the pollutant, the greater the incentive to wait rather than adopt the policy now, and hence the greater must be the current cost in order to trigger adoption. Second, an increase in the discount rate r increases the value of the option to adopt the policy and thus also increases θ^* . This is analogous to the effect of a change in the interest rate on the value and optimal exercise point for a financial option; an increase in r implies that lead to a bigger reduction in the present value of the cost K of policy adoption, so that the option to adopt is worth more but it should be exercised later. Third, an increase in δ , the rate of “depreciation” of the stock of pollutant, also increases θ^* ; a higher δ implies that the environmental damage from emissions is less irreversible, so that the sunk benefit of adopting the policy now rather than waiting is reduced.

Also, observe that an increase in the initial rate of emissions E_0 leaves θ^* unchanged (but increases the value of society’s option to adopt the emission-reducing policy). the reason is that $K = kE_0$, so that θ^* is independent of E_0 , and A increases linearly with E_0 . Finally, θ^* is also independent of M . Because $B(M, \theta)$ is linear in M (so that the value function is linear in M), any given level of M_t implies the same reduction in social welfare if the policy is adopted at time t as it does if the policy is not adopted. Hence $W^N(\theta, M) - W^A(\theta, M)$ is independent of M , and so is θ^* .

In Section 3 we saw that the policy timing problem could be framed in terms of a comparison of the opportunity costs of current adoption with the corresponding opportunity “benefits.” We can do the same thing here. In Section 3 we calculated $W_T - W_0$, i.e., the value function when a decision is made at a (fixed) future time T less the value function when

the policy is adopted immediately. In the current model, T is unknown, but we can calculate $W^* - W_0$, where W^* is the value function when the adoption decision is made optimally. Suppose $\theta < \theta^*$, so that $W^* = W^N$. Then $W^* - W_0 = W^N - W_0 = W^N - W^A + K$, or

$$W^* - W_0 = K + A\theta^\gamma - \frac{\beta E_0 \theta}{(r - \alpha)(r + \delta - \alpha)}. \quad (49)$$

The first term on the right-hand side of (49) is the direct cost of current adoption. The second term is value of the option to adopt, and since adoption implies “killing” this option, it is an opportunity cost of current adoption. The last term is the present value of the additional flow of social cost from continued emissions, and thus is an opportunity “benefit” of current adoption. Since $\theta < \theta^*$ and $W^* - W_0 < 0$, the direct cost and opportunity cost outweigh this opportunity benefit, and adoption should be delayed.

Note that the way this model is currently structured, we would never want to reduce emissions by anything less than 100 percent (assuming we would want to reduce them at all). The reason is that with $K = kE_0$, the value of the option to adopt the policy, $A\theta^\gamma$, is linear in E_0 , so that W^N and W^A are linear in M and E_0 . Shortly we will generalize the model so that K is nonlinear. Then we can examine policies that involve a one-time partial reduction in emissions, as well as gradual incremental reductions in emissions.

As before, we will use a numerical example to explore the characteristics of the solution. Suppose $\alpha = 0$ (so that the expected social cost per unit of M is constant), $r = .04$, $\delta = .02$, $\sigma = .20$, $\beta = 1$, $E_0 = 300,000$ tons per year, $\theta_0 = \$20$ per ton, and $k = 6667$ so that $K = \$2$ billion. Then $\gamma = 2.0$, $A = 1,953,125$, and $\theta^* = \$32$ per ton. Hence at the current value of $\theta_0 = 20$, the policy should not be adopted, but the value of society’s option to adopt it, $A\theta^\gamma$, is \$0.78 billion. The policy should only be adopted when θ reaches \$32 per ton; at that point $A\theta^\gamma = \$2.0$ billion, and the reader can check that boundary conditions (42) and (43) are satisfied.

Figure 4 shows this solution graphically for $M = 0$ (so that $W^A = 0$ for all values of θ). Note that θ^* is found at the point of tangency of W^N with the line $W^A - K$. (If M were greater than zero, we would have $W^A = -\theta M / (r + \delta - \alpha)$, so we would simply rotate both $W^N(\theta)$ and the line $W^A - K$ downwards.) Figure 5 shows how the solution in Figure 4




Figure 4: Solution for $M = 0$.




Figure 5: Effect of Increasing σ on θ^* .




Figure 6: θ^* as a Function of σ , $\delta = .01, .02$.

changes when σ is increased. It shows the value functions and points of tangency for σ equal to .2 and .4, with the other parameters held constant. Note that when σ is increased from .2 to .4, γ falls (from 2 to 1.366), so that $A\theta^\gamma$ becomes less convex, $W^N(\theta)$ flattens out, and θ^* approximately doubles to a value of 59.72.

Figure 6 shows θ^* as a function of σ for two different values of δ (.01 and .02), and Figure 7 shows θ^* as a function of δ for $\sigma = 0.2$ and 0.4. Note that for all values of δ , θ^* rises sharply with σ . This is partly due to the fact that we have framed the policy problem as an all-or-nothing proposition, but it nonetheless suggests that assessing uncertainty over the future costs and benefits of emission reduction may be particularly critical to the policy adoption decision.

Convex Costs and Partial Reduction in Emissions.

In Section 3 we examined policies that would only partially reduce emissions, but in the context of a two-period decision framework (i.e., adopt the policy now or wait until a fixed time T to decide). We now return to that problem, but allowing the policy to be adopted at any time. As before, we will assume that the cost of the policy is a quadratic function of




Figure 7: θ^* as a Function of δ , $\sigma = .2, .4$.

the amount that emissions are reduced:

$$K = k_1(E_0 - E_1) + k_2(E_0 - E_1)^2, \quad (50)$$

where $E_0 - E_1$ is the amount of the reduction, and $k_1, k_2 > 0$. Again, the cost of a 1-unit (permanent) reduction in E is $k(E) = -dK/dE_1 = k_1 + 2k_2(E_0 - E_1)$. As before, we must find a rule (in the form of a critical value θ^*) for the optimal timing of policy adoption, but now we must also determine the optimal size of the reduction.

Once again, let $W^N(\theta, M)$ and $W^A(\theta, M)$ be the value functions for the “no-adopt” and “adopt” regions respectively. $W^N(\theta, M)$ again satisfies eqn. (39). However, eqn. (40) for $W^A(\theta, M)$ now has an additional term on the right-hand side: $\beta E_1 W_M^A$, where E_1 is the level of emissions after the policy has been adopted. Hence the solution for $W^A(\theta, M)$ is now:

$$W^A(\theta, M) = -\frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_1 \theta}{(r - \alpha)(r + \delta - \alpha)}, \quad (51)$$

while the solution for $W^N(\theta, M)$ is again given by eqn. (44).

Remember that E_1 is chosen optimally, and so will depend on θ at the time of adoption, i.e., on θ^* . Hence although boundary condition (41) will still apply, conditions (42) and (43)

must be rewritten as:

$$W^N(\theta^*, M) = W^A(\theta^*, M) - K(E^*(\theta^*)), \quad (52)$$

and

$$W_\theta^N(\theta^*, M) = W_\theta^A(\theta^*, M) - \frac{dK}{dE^*} \frac{dE^*}{d\theta^*}. \quad (53)$$

Using eqns. (51) and (50), we choose E^* to maximize the net payoff from policy adoption:

$$\begin{aligned} \max_E [W^A(\theta, M; E) - K(E)] = & - \frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_1 \theta}{(r - \alpha)(r + \delta - \alpha)} \\ & - k_1(E_0 - E) - k_2(E_0 - E)^2, \end{aligned} \quad (54)$$

so that

$$E^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta \theta}{2k_2(r - \alpha)(r + \delta - \alpha)}. \quad (55)$$

We now substitute this expression for E^* into boundary conditions (52) and (53), and then use these conditions to find θ^* and A . Making the substitutions and denoting $\rho \equiv (r - \alpha)(r + \delta - \alpha)$, we find that θ^* is the largest root¹⁵ of the quadratic equation $(\gamma - 2)\beta^2\theta^2 - 2\rho(\gamma - 1)\beta k_1\theta + \gamma\rho^2 k_1^2 = 0$, i.e.,

$$\theta^* = \frac{\rho(\gamma - 1)k_1}{\beta(\gamma - 2)} \left[1 + \sqrt{1 - \frac{\gamma(\gamma - 2)}{(\gamma - 1)^2}} \right], \quad (56)$$

and A is given by:

$$A = \beta^2/4k_2\rho^2(\theta^*)^{\gamma-2} - \beta k_1/2k_2\rho(\theta^*)^{\gamma-1} + k_1^2/4k_2(\theta^*)^\gamma. \quad (57)$$

Given θ^* , we can find $E^*(\theta^*)$ from eqn. (55). It is easy to confirm that as σ increases (so that γ decreases), θ^* will increase and E^* will fall. However, we must have $E^* \geq 0$, which implies $\theta^* \leq \theta_m = \rho k_1/\beta + 2\rho k_2 E_0/\beta$, or equivalently, $\gamma \geq 2 + k_1/k_2 E_0$.¹⁶ If σ is large enough so that $\gamma < 2 + k_1/k_2 E_0$, eqns. (56) and (57) will no longer apply. Instead, E^* is constrained to be zero and so is no longer a choice variable. In that case the solution to the optimal

¹⁵This is because $W^A(\theta) - W^N(\theta) - K(E^*(\theta))$ is convex in θ .

¹⁶Thus the condition from eqn. (56) that $\gamma > 2$, which implies that $\sigma^2 < r - 2\alpha$, will always be satisfied.




Figure 8: Optimal Adoption Decision with Partial Emission Reduction.

timing problem is again given by eqns. (47) and (48) (with $K = k_1 E_0 + k_2 E_0^2$). Also, we must have $E^* \leq E_0$, but this will always be the case; observe from eqns. (55) and (56) that $E^*(\theta^*) < E_0$ for any $\gamma > 2$.¹⁷

Figure 8 illustrates this. It shows the conventional NPV from policy adoption when E_1 can be chosen optimally according to eqn. (55); that NPV is equal to $(\beta/\rho)[E_0 - E^*(\theta)]\theta - K(E^*(\theta))$. Also shown is the value of the option to adopt a policy, which is equal to $A\theta^\gamma$. This NPV is a quadratic function of θ , and applies for values of θ for which $0 \leq E^*(\theta) \leq E_0$; in this range, the NPV is increasing in θ . The critical value θ^* is at the point where the option value $A\theta^\gamma$ is just tangent to this NPV, i.e., where the value matching and smooth pasting conditions (52) and (53) hold. In the figure, that critical value is shown as θ_1^* . If

¹⁷It might appear from eqn. (55) that if β is very small, E^* will exceed E_0 . But as β becomes smaller, θ^* becomes larger, so that $E^* < E_0$ always.




Figure 9: Partial Emission Reduction — $k_1 = 5000$, $k_2 = .055$, $\sigma = .045$.

σ is large enough (or other parameters are sufficiently large or small) that $\theta^* > \theta_{\max}$, E_1 is constrained to be zero, so that the corresponding NPV is equal to $(\beta/\rho)E_0\theta - k_1E_0 - k_2E_0^2$. In the figure, a solution in this range is indicated by θ^* .

Figure 9 shows this solution for the following numerical example: $E_0 = 300,000$ tons per year, $k_1 = 5000$ and $k_2 = .0055$ (so that the cost of reducing E to zero would be about \$2 billion), $\sigma = .045$, and as before, $\alpha = 0$, $r = .04$, $\delta = .02$, and $\beta = 1$. In this case, a policy is never adopted for $\theta < \theta_{\min} = 12$ (even if σ is reduced to zero), and $\theta_{\max} = 20$. For $\sigma = .045$, $\gamma = 6.8$, so that $\theta^* = 17$, i.e., $\theta^* < \theta_{\max}$ so that $E^* > 0$. From eqn. (55)), we see that $E^* = 110,606$ tons per year.

The amount that emissions are reduced will depend on the degree of uncertainty over the future benefits of reduction, and on other parameters. Figure 10 shows the dependence of both E^* and θ^* on σ for this numerical example. (In the figure, θ^* is multiplied by 10^4 so that it can be plotted with E^* on the same scale.) When σ is equal to zero, the standard NPV rule for the timing of adoption applies; the policy should be adopted if $\theta \geq 12$. If θ is just slightly greater than 12, the policy is adopted but emissions are reduced only very




Figure 10: Partial Emission Reduction — Dependence of E^* and θ^* on σ .

slightly. (The reason is that $\alpha = 0$, so if $\sigma = 0$, θ cannot rise in the future.) As σ increases, the critical value θ^* also increases, and E^* falls. Note that for $\sigma > .063$, $E^* = 0$, so that θ^* is given by eqn. (48) (with $k_1 E_0 + k_2 E_0^2$ substituted for K) rather than eqn. (56).

We can likewise determine the dependence of θ^* and E^* on other parameters from eqns. (48), (55), and (56). For example, a higher initial level of emissions, E_0 , does not affect the critical value θ^* , but does imply a commensurately higher ending level E^* (so that the size of the reduction is unchanged). Also, an increase in k_1 increases θ^* , but an increase in k_2 has no effect on θ^* , although it reduces the size of the emission reduction, $E_0 - E^*$, at the time of adoption.

Convex Benefit Function.

In all of the cases examined so far, we have assumed that the benefit function $B(M, \theta)$ is linear in M , and as a result, the optimal policy rules were always independent of M . This was a convenient assumption that made it easier to analytically solve for the optimal policy rule. However, for most environmental problems, the damage from a pollutant is like to rise more than proportionally with the stock of the pollutant. If this is the case, the

optimal policy rule (i.e., the optimal timing and amount of emission reduction) will depend on the stock, M . To explore this, let us go back to the case in which the cost of an emission reduction is linear in the size of the reduction, and emissions must be reduced to zero once a policy is adopted, so that $K = kE_0$. But now we will assume that the benefit function $B(M, \theta)$ is quadratic in M :

$$B(M, \theta) = -\theta_t M_t^2. \quad (58)$$

Once again, we can write value functions $W^N(\theta, M)$ and $W^A(\theta, M)$ for the “no-adopt” and “adopt” regions respectively. In this case, these value functions will again satisfy the Bellman equations (39) and (40), but with the term $-\theta M$ replaced by $-\theta M^2$ in each equation. The boundary conditions (41) – (43) also apply. These equations and boundary conditions have the following solution:

$$\begin{aligned} W^N(\theta, M) = A\theta^\gamma & - \frac{\theta M^2}{r + 2\delta - \alpha} - \frac{2\beta^2 E_0^2 \theta}{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)} \\ & - \frac{2\beta E_0 \theta M}{(r + 2\delta - \alpha)(r + \delta - \alpha)}, \end{aligned} \quad (59)$$

and

$$W^A(\theta, M) = -\frac{\theta M^2}{r + 2\delta - \alpha}, \quad (60)$$

where A is a positive constant to be determined, and γ is again given by eqn. (eq:thlingamma). Note that the right-hand side of (60) and the second term on the right-hand side of (59) is the present value of the flow of social cost from the present stock of pollutant, M . The third and fourth terms on the right-hand side of (59) are the present value of the flow of social cost from future emissions at the rate E_0 , and the first term on the right-hand side of (59) is the value of the option to reduce emissions to zero.

The constant A and critical value θ^* , which are found from boundary conditions (42) and (43), are given by:

$$A = E \left(\frac{\gamma - 1}{k} \right)^{\gamma-1} \left[\frac{2\beta^2 E_0 + 2\beta(r - \alpha)M}{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)\gamma} \right]^\gamma, \quad (61)$$

$$\theta^* = \frac{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)k\gamma}{2\beta(\gamma - 1)[\beta E + (r - \alpha)M]}. \quad (62)$$

The critical value θ^* now depends on M ; a higher value of M implies a higher marginal social cost from additional emissions, and therefore a lower value of θ at which it is optimal to begin reducing emissions. (For the same reason, a higher value of M increases the value of the option to reduce emissions.) The rising marginal social cost of emissions likewise implies that the higher is the current emission level, E_0 , the lower will be θ^* . As before, a higher cost of emission reduction, k , and a higher rate of “depreciation” of the stock of pollutant, δ , leads to a higher value of θ^* .

Most important, uncertainty affects the optimal adoption rule the same way it does when $B(\theta, M)$ is linear in M ; the parameter σ affects θ^* through the multiplier $(\gamma - 1)/\gamma$, and γ is given by the same equation (46) as before. Hence we find that making the benefit function convex in M affects the optimal policy adoption rule, but it does not affect the way that rule depends on uncertainty over the future social costs of pollution. The critical value θ^* for the certainty case is multiplied by the same factor as it is when the benefit function is linear in M .

6. Cost-Benefit Uncertainty — Gradual Emission Reductions.

In the preceding section we assumed that there would be only one opportunity to adopt an emissions-reducing policy. This is not terribly unrealistic; given the political difficulties of reaching a consensus on and introducing a major new environmental policy, it is unlikely that regulations regarding emissions could be revised very frequently. On the other hand, assuming that such regulations could never be revised (once a new policy is in place) is extreme. Rather than making arbitrary assumptions about the allowed frequency of policy change (or making assumptions about “menu costs” of policy change so that the frequency is endogenous), I will assume the opposite extreme — that the level of emissions can be reduced gradually and continuously. Comparing the optimal policy in this case with that from the preceding section should provide some insight into how the frequency with which regulations can be introduced or substantively changed will affect the optimal timing and design of policy.

In this section I will again assume that the cost of any incremental emission reductions is

completely sunk, which is equivalent to assuming that emissions can only be reduced. (This assumption can easily be relaxed by making the cost of emission reductions only partly sunk.) Policy makers must observe the stock variable M , and decide when and by how much to mandate emissions reductions in response to increases in M .¹⁸

For this problem to be of any interest, either the benefit function or the cost function must be convex. I will assume that the benefit function $B(\theta, M)$ is again linear in θ and M , and that the cost of the policy is a quadratic function of the amount that emissions are reduced, as in eqn. (50). Thus the cost of a 1-unit reduction in E is $\Delta K = k_1 + 2k_2(E_0 - E_1)$. For notational convenience, let $m_1 = k_1 + 2k_2E_0$ and $m_2 = 2k_2$, so that the cost of an incremental reduction in E can be written as:

$$\Delta K = m_1 - m_2 E. \quad (63)$$

Since $B_t = -\theta_t M_t$, the payoff flow from a small reduction in the stock of pollutant, ΔM_t , is just $\Delta B_t = -\theta_t \Delta M_t$. If emissions are reduced incrementally by an amount ΔE at time $t = 0$, the corresponding change in M_t is

$$\Delta M_t = -\frac{\beta \Delta E}{\delta} [1 - e^{-\delta t}], \quad (64)$$

so the social benefit from an incremental reduction in emissions at time t is:

$$\Delta W_t = \mathcal{E}_t \int_t^\infty \Delta B_\tau e^{-r(\tau-t)} d\tau = \beta \theta_t \Delta E / \rho, \quad (65)$$

where $\rho \equiv (r - \alpha)(r + \delta - \alpha)$. Given the current value of θ_t , the problem is to determine how far to reduce emissions initially, and then how to make any further reductions in response to changes in θ .

This problem is analogous to the incremental investment and capacity choice problem in Pindyck (1988). (See also Dixit and Pindyck (1994), Chapter 11.) Suppose that $E_t = E$

¹⁸I am implicitly assuming here that M can be observed without error. In many cases — e.g., GHG concentrations in the atmosphere — one would only have an estimate of M . Later in this paper I examine how uncertainty over the *evolution* of M should affect the policy decision, but I do not examine the implications of uncertainty over the *current value* of M . Also, note that in some cases, one could, through the expenditure of time and money, make it possible to obtain better estimates of M . (This has been one of the goals of recent geophysical research on global warming.) This leads to a stochastic control problem in which resources spent on estimation must be balanced against resources spent on direct control of emissions. See Rausser and Howitt (1975) for an example.

currently, and let $W(E; \theta, M)$ be the value function given this E , and given the current values of θ and M . Let ΔF be the value of society's option to (permanently) reduce E by one unit. Note that the cost of exercising that option is $\Delta F(E; \theta, M) + \Delta K(E)$, and the payoff is $\Delta W(\theta)$. Then ΔF must satisfy the Bellman equation:

$$r\Delta F = (\beta E - \delta M)\Delta F_M + \alpha\theta\Delta F_\theta + \frac{1}{2}\sigma^2\theta^2\Delta F_{\theta\theta}, \quad (66)$$

subject to the boundary conditions:

$$\Delta F(E; 0, M) = 0, \quad (67)$$

$$\Delta F(E; \theta^*, M) = \Delta W(\theta^*) - \Delta K(E), \quad (68)$$

$$\Delta F_\theta(E; \theta^*, M) = \Delta W_\theta(\theta^*). \quad (69)$$

Since ΔW and ΔK are independent of M , ΔF will be independent of M , and the solution has the usual form:

$$\Delta F = a\theta^\gamma, \quad (70)$$

with $\gamma > 1$ again given by eqn. (46). Emissions should be reduced whenever θ exceeds the critical value $\theta^*(E)$, with $d\theta^*/dE < 0$. The constant a and the critical value $\theta^*(E)$ are found from boundary conditions (68) and (69):

$$\theta^*(E) = \frac{\gamma\rho(m_1 - m_2E)}{(\gamma - 1)\beta}, \quad (71)$$

$$a = \left(\frac{\beta}{\gamma\rho}\right)^\gamma \left(\frac{\gamma - 1}{m_1 - m_2E}\right)^{\gamma-1}, \quad (72)$$

where $\rho \equiv (r - \alpha)(r + \delta - \alpha)$.

To interpret (71), note that $\rho(m_1 - m_2E)/\beta$ is the amortized sunk cost of an incremental reduction in emissions, normalized by the absorption rate β . Since the benefit function $B(\theta, M)$ is linear, in the absence of uncertainty it would be optimal to reduce emissions to the point where this amortized sunk cost is just equal to θ , the social cost per period of an incremental unit of the stock of pollutant, M . With uncertainty, the threshold exceeds this amortized sunk cost by the multiple $\gamma/(\gamma - 1)$. Also, note that as E is reduced, θ^*

rises (and a falls). Depending on the initial value of θ , it may be optimal to initially reduce emissions by some large amount, and then later reduce emissions gradually when θ increases and hits the boundary θ^* . For any value of E , θ^* is increased if σ increases, and is decreased if the “depreciation” rate δ increases. Finally, given $\theta^*(E)$, we can determine the optimal emissions level E^* .

In this model, uncertainty affects the initial level of mandated emissions reductions, and it also affects the maximum allowed emissions level over time. I used a numerical example and ran a Monte Carlo simulation to examine the magnitude of these effects and its dependence on σ . In this example, the initial emissions level is $E_0 = 300,000$ tons per year, the cost function parameters for eqn. (63) are $k_1 = 5000$ and $k_2 = .0055$ (so that the cost of reducing E from 300,000 tons per year to zero would be about \$2 billion), and as in earlier examples, $r = .04$, $\delta = .02$, and $\beta = 1$. In this example, we set $\alpha = .01$, so that even in the absence of uncertainty, emissions will gradually be reduced over time as θ increases. I varied σ from 0 to .15, in increments of .005. For each value of σ , I ran 10,000 simulations of the evolution of θ and the corresponding optimal emissions level E^* .

Figure 11 shows the results of this Monte Carlo simulation for the mean optimal emissions level initially, and after 20 years. Note that when there is no uncertainty (i.e., $\sigma = 0$), emissions are initially reduced from 300,000 to about 70,000 tons per year, and then reduced gradually to zero as θ and the corresponding social cost of pollution rises. As σ is increased, the initial allowed emissions level increases, reflecting the value of waiting. Emissions are still reduced over time (although reductions occur stochastically when $\sigma > 0$), but the mean value of E^* after 20 years also increases with σ .

Figure 12 shows the mean and median times until the optimal emissions level has been reduced to zero. Both the mean and median times should increase monotonically with σ , because increases in σ increase the threshold $\theta^*(E)$ for every value of E . In the figure, the mean time decreases for $\sigma > .13$, but this is an artifact of the Monte Carlo simulation. (In each run, the model was simulated for 1000 years, and for large values of σ , there will be runs for which it takes longer than this for E^* to reach zero. In addition, the number of runs at this tail of the distribution is very small.) Note that because the distribution of the




Figure 11: Mean Optimal Emissions Level at $t = 0$ and 20 Years.




Figure 12: Mean and Median Times until Emissions are Reduced to Zero.

time until zero emissions is asymmetric, the mean time will exceed the median time for all $\sigma > 0$. The difference between the mean time and median time illustrates an important aspect of the effects of uncertainty. There is a value of waiting (i.e., reducing emissions less than would be the case otherwise) because of the possibility that θ will not increase as much as expected. For $\sigma > 0$, there are indeed realizations in which it takes a very long time for θ to grow to the point where eliminating emissions is justified. This is another example of the “good news principle” mentioned earlier in this paper — the tail of the distribution drives the initial amount by which emissions are reduced.

In this section and the preceding one, we considered two extreme assumptions regarding policy adoption. In the preceding section we assumed that there would be only one opportunity to adopt an emissions-reducing policy, whereas in this section we assumed that emissions could be reduced continuously over time. Reality lies somewhere in between these two extremes. Nonetheless, the basic implications of uncertainty are the same — there is an incentive to go slow, delaying policy adoption in the one-time-only case, and initially reducing emissions by a smaller amount in the continuous case.

7. Ecological Uncertainty.

Up to this point, the only form of uncertainty that we have considered has been over the parameter θ that shifts the benefit function B . In this section, we will assume that θ remains fixed, but that there is uncertainty over the evolution of M . Specifically, we replace eqn. (1) by the following stochastic differential equation:

$$dM = (\beta E - \delta M)dt + \sigma dz. \tag{73}$$

Thus even if the trajectory for E_t were known, future values of M are uncertain (and normally distributed).¹⁹

¹⁹It might seem more natural to assume that future values of M are lognormally distributed, i.e., to describe the evolution of M by

$$dM = (\beta E - \delta M)dt + \sigma M dz.$$

Then M could never become negative. I use eqn. (73) instead because it simplifies the numerical solution of the model. The basic results would still apply if M were lognormally distributed.

For uncertainty of this kind to have any effect on policy timing or design, the benefit function $B(\theta, M)$ must be convex in M . (If this function were linear in M , stochastic fluctuations in M would have no effect on the expected marginal social return from reductions in E , and thus would not affect the optimal policy, even if $K(E)$ were nonlinear.) We will therefore assume that the benefit function is quadratic in M , i.e., $B(M, \theta) = -\theta M_t^2$. For simplicity, we will also assume that the cost of an emission reduction is linear in the size of the reduction, and emissions must be reduced to zero once a policy is adopted, so that $K = kE_0$.

We can now proceed as before, writing the Bellman equations for the value functions W^N and W^A in the “no-adopt” and “adopt” regions:

$$rW^N = -\theta M^2 + (\beta E_0 - \delta M)W_M^N + \frac{1}{2}\sigma^2 W_{MM}^N, \quad (74)$$

$$rW^A = -\theta M^2 - \delta M W_M^A + \frac{1}{2}\sigma^2 W_{MM}^A. \quad (75)$$

The value functions must also satisfy the boundary conditions:

$$W^A(0) = 0, \quad (76)$$

$$W^N(M^*) = W^A(M^*) - K, \quad (77)$$

$$W_M^N(M^*) = W_M^A(M^*), \quad (78)$$

where M^* is the critical value of M that triggers policy adoption.

Note that there is now only one state variable (M), so that eqns. (74) and (75) are ordinary differential equations. The solution for W^A is simply:

$$W^A(M) = -\frac{\theta M^2}{r + 2\delta} - \frac{\sigma^2 \theta}{r(r + 2\delta)}. \quad (79)$$

This is just the present value of the flow of social cost from the current stock of the pollutant, M . Note that an increase in σ implies an increase in the magnitude of W^A . This is an implication of Jensen’s inequality; W^A is a convex function of M .

Eqn. (74) for W^N , however, does not have an analytical solution. Instead, this equation must be solved numerically for $W^N(M)$ and the critical value M^* , by making use of the




Figure 13: Solution Method for Stochastic M .
 (Note: $K = 4$, $E_0 = .3$, $\sigma = 1$, $\theta = .002$.)

solution for $W^A(M)$ given by eqn. (79), along with the boundary conditions (77) and (78). Obtaining a numerical solution in this case is straightforward. Begin with a candidate for the critical value M^* that triggers policy adoption, say, M_0^* . Then use eqns. (79), (77), and (78) to get $W^N(M_0^*)$ and $W_M^N(M_0^*)$, and solve eqn. (74) backwards to determine a candidate solution for $W^N(M)$ for all M between 0 and M_0^* . To be the actual solution, the candidate solution must, for all values of M between 0 and M_0^* , satisfy the conditions $W_M^N < 0$ and $W_{MM}^N < 0$. The candidate for M^* is adjusted until these conditions are satisfied.

This solution method is easiest to see in the context of an example. We will measure the stock of pollutant, M , in millions of tons, the emission rate in millions of tons per year, and the value functions W^A and W^N and adoption cost in billions of dollars. Since the benefit function is $B = -\theta M^2$, we measure θ in billion dollars/(million tons)². Then, we set $K = 4$, $E_0 = .3$, $\theta = .002$, $\sigma = 1$, $\alpha = 0$, and, as in our earlier examples, $r = .04$, $\delta = .02$, and $\beta = 1$. In this case, the solution for M^* is 13.05. Figure 13 shows how this solution is obtained. The figure shows candidate solutions for $W^N(M)$ corresponding to different values of M^* , along




Figure 14: Solution for Stochastic M — $K = 4$, $E_0 = .3$, $\sigma = 1$, $\theta = .002$.

with $W^A(M) - K$. Note that for candidate values of M^* below 13.05, $W_M^N(M) > 0$ for small values of M , and for candidate values above 13.05, $W_{MM}^N(M) > 0$ for small values of M . The solution procedure simply searches over candidate values of M^* , using an increasingly narrow range.

Figure 14 shows this same solution, but for values of M ranging from 0 to 20 million tons. Note that for $M < M^*$, the curve for $W^N(M)$ applies, and for $M > M^*$, $W^A(M)$ applies. Finally, Figure 15 shows how this solution changes in response to changes in the value of σ . Solutions are shown for $\sigma = 0, 1$, and 2. Observe that as σ is increased, the critical value M^* increases. This is a standard result; the value of the option to adopt the policy increases, and hence the value of waiting increases. Likewise, the expected present value of the social cost of pollution, before or after policy adoption, becomes larger.

8. A General Model.

In the general case, both θ and M are stochastic. We will assume as we did in Section 5 that θ follows geometric Brownian motion:

$$d\theta = \alpha\theta dt + \sigma_1\theta dz_1, \quad (80)$$

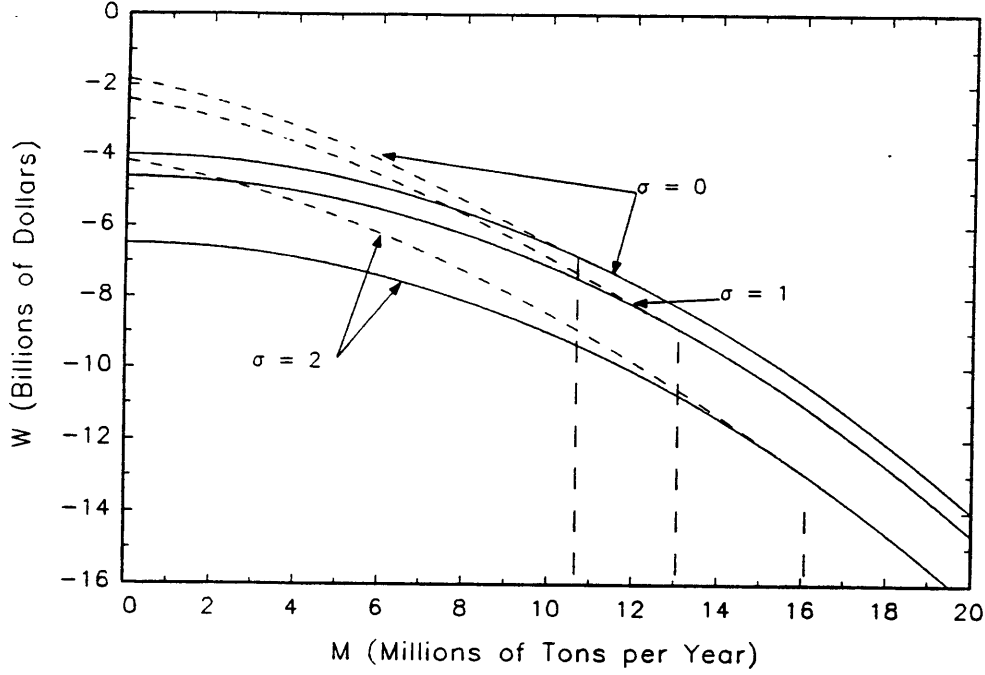


Figure 15: Solution for Stochastic M — $\sigma = 1, 2,$ and 3 .

and that M follows a controlled arithmetic Brownian motion:

$$dM = (\beta E - \delta M)dt + \sigma_2 dz_2, \quad (81)$$

with $\mathcal{E}(dz_1 dz_2) = 0$. We will also assume that the social benefit function is quadratic, i.e., $B(\theta, M) = -\theta M^2$.

We proceed as before, writing the Bellman equations for the value functions $W^N(\theta, M)$ and $W^A(\theta, M)$ in the “no-adopt” and “adopt” regions respectively:

$$rW^N = -\theta M^2 + (\beta E_0 - \delta M)W_M^N + \alpha \theta W_\theta^N + \frac{1}{2}\sigma_1^2 \theta^2 W_{\theta\theta}^N + \frac{1}{2}\sigma_2^2 W_{MM}^N, \quad (82)$$

$$rW^A = -\theta M^2 - \delta M W_M^A + \alpha \theta W_\theta^A + \frac{1}{2}\sigma_1^2 \theta^2 W_{\theta\theta}^A + \frac{1}{2}\sigma_2^2 W_{MM}^A. \quad (83)$$

These value functions must also satisfy the boundary conditions:

$$W^A(0, M) = 0, \quad (84)$$

$$W^N(0, M) = 0, \quad (85)$$

$$W^N(\theta^*(M), M) = W^A(\theta^*(M), M) - K, \quad (86)$$




Figure 16: Free Boundary for General Model.

$$W_{\theta}^N(\theta^*(M), M) = W_{\theta}^A(\theta^*(M), M). \quad (87)$$

Now the free boundary that separates the adopt from the no-adopt regions is a function $\theta^*(M)$, with $d\theta^*/dM < 0$. Given M , the policy should be adopted only if $\theta > \theta^*(M)$. Figure 16 illustrates this.

As in the case when only M was stochastic, we can obtain an analytical solution for $W^A(\theta, M)$:

$$W^A(\theta, M) = -\frac{\theta M^2}{r - \alpha + 2\delta} - \frac{\sigma_2^2 \theta}{(r - \alpha)(r - \alpha + 2\delta)}. \quad (88)$$

As before, this is the present value of the flow of social cost from the current stock of the pollutant, M . Again, an increase in σ implies an increase in the magnitude of W^A , because W^A is a convex function of M .

Eqn. (82) for $W^N(\theta, M)$ does not have an analytical solution. However, it can be solved numerically for $W^N(\theta, M)$ and for the free boundary $\theta^*(M)$, by making use of the solution for $W^A(\theta, M)$ along with the boundary conditions (86) and (87).

References

- Arrow, Kenneth J., and Anthony C. Fisher. "Environmental Preservation, Uncertainty, and Irreversibility," *Quarterly Journal of Economics*, 1974, **88**, 312–319.
- Bernanke, Ben S., "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, Feb. 1983, **98**, 85–106.
- Cline, William R., *The Economics of Global Warming*, Institute for International Economics, Washington, D.C., 1992.
- Dixit, Avinash, "Investment and Hysteresis," *Journal of Economic Perspectives*, Winter 1992, **6**, 107–132.
- Dixit, Avinash and Robert S. Pindyck, *Investment Under Uncertainty*, Princeton, N.J.: Princeton University Press, 1994.
- Dixit, Avinash and Robert S. Pindyck, "A Generalized Option-Theoretic Model of Investment," draft working paper, March 1995.
- Fisher, Anthony C., and W. Michael Hanemann, "Information and the Dynamics of Environmental Protection: The Concept of the Critical Period," *Scandinavian Journal of Economics*, 1990, **92**, 399–414.
- Hammitt, James K., Robert J. Lempert, and Michael E. Schlesinger, "A Sequential-Decision Strategy for Abating Climate Change," *Nature*, May 28, 1992, **357**, 315–318.
- Hanemann, W. Michael, "Information and the Concept of Option Value," *Journal of Environmental Economics and Management*, Jan. 1989, **16**, 23–37.
- Hendricks, Darryll, "Optimal Policy Response to an Uncertain Threat: The Case of Global Warming," unpublished manuscript, Kennedy School of Government, Harvard University, March 1992.
- Henry, Claude, "Investment Decisions under Uncertainty: The Irreversibility Effect," *American Economic Review*, Dec. 1974, **64**, 1006–12.
- Jacoby, Henry D., and Ronald G. Prinn, "Uncertainty in Climate Change Policy Analysis," Report No. 1, MIT Joint Program on the Science and Policy of Global Change, December 1994.
- Kolstad, Charles D., "Regulating a Stock Externality Under Uncertainty with Learning,"

- Working Paper No. 92-0112, Department of Economics, University of Illinois at Urbana-Champaign, March 1992.
- Kolstad, Charles D., "Learning and Irreversibilities in Environmental Regulation: The Case of Greenhouse Gas Emissions," Working Paper No. 94-001, Department of Economics, University of Illinois at Urbana-Champaign, Jan. 1994.
- Krutilla, John V., and Anthony C. Fisher, *The Economics of Natural Environments*, Baltimore: Johns Hopkins University Press, 1975.
- Lind, Robert C., *Discounting for Time and Risk in Energy Policy*, Baltimore: Johns Hopkins University Press, 1984.
- Marglin, Stephen, *Approaches to Dynamic Investment Planning*, Amsterdam: North-Holland, 1963.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw, "The Impact of Global Warming on Agriculture: A Ricardian Analysis," *American Economic Review*, September 1994, **84**, 753-771.
- Nordhaus, William D., "To Slow or Not to Slow: The Economics of the Greenhouse Effect," *The Economic Journal*, July 1991, **101**, 920-937.
- Pindyck, Robert S., "Irreversible Investment, Capacity Choice, and the Value of the Firm," *American Economic Review*, Dec. 1988, **79**, 969-985.
- Pindyck, Robert S., "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, Sept. 1991, **29**, 1110-1152.
- Rausser, Gordon C., and Richard Howitt, "Stochastic Control of Environmental Externalities," *Annals of Economic and Social Measurement*, Spring 1975, **4**, 271-292.
- Schmalensee, Richard, Thomas M. Stoker, and Ruth A. Judson, "World Energy Consumption and Carbon Dioxide Emissions: 1950-2050," MIT Center for Energy and Environmental Policy Research, Working Paper 95-001WP, February 1995.
- Solow, Andrew R., "Is There a Global Warming Problem?" in R. Dornbusch and J. Poterba, eds., *Global Warming: Economic Policy Responses*, Cambridge, MA: MIT Press, 1991.
- Weitzman, Martin, "The Environmental Discount Rate," *Journal of Environmental Economics and Management*, 1994.