

Irreversibilities and the timing of environmental policy

Robert S. Pindyck*

*Sloan School of Management, E52-453, Massachusetts Institute of Technology,
Cambridge, MA 02142, USA*

Received 12 January 1999; accepted 24 May 1999

Abstract

The standard framework in which economists evaluate environmental policies is cost–benefit analysis, so policy debates usually focus on the expected flows of costs and benefits, or on the choice of discount rate. But this 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 two kinds of uncertainty — over the future costs and benefits of reduced environmental degradation, and 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. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: Q28; L51; H23

Keywords: Environmental policy; Irreversibilities; Cost–benefit analysis; Uncertainty; Option value; Global warming

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.

* Tel.: +1-617-253-6641; fax: +1-617-258-6855.

E-mail address: rpindyck@mit.edu (R.S. Pindyck).

Presumably, it also yields an expected flow of benefits. 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.

This standard framework ignores three important characteristics of most environmental problems and the policies designed to respond to them. First, there is almost always uncertainty over the future costs and benefits of adopting a particular policy. With global warming, for example, we do not know how much average temperatures will rise with or without reduced CO₂ emissions, nor do we know 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 adopting policies to reduce the damage. Third, policy adoption is rarely a now-or-never proposition. In most cases, it is feasible to delay action and wait for new information, or at least begin with policies that are limited in their scope and impact.

Environmental policy involves 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.

Second, environmental damage can be partially or totally irreversible. For example, increases in greenhouse gas (GHG) concentrations are long lasting; even if radical policies were adopted to drastically reduce GHG emissions, these concentrations would take many years to fall. Also, the damage to ecosystems from higher global temperatures (or from acidified lakes and streams, or 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, which 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.¹

¹ This point was made some two decades ago by Arrow and Fisher (1974), Henry (1974), and Krutilla and Fisher (1975), and has been elaborated upon by Fisher and Hanemann (1990), among others.

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 types of uncertainty are relevant. The first is *economic uncertainty*, i.e., uncertainty over the future costs and benefits of environmental damage and its reduction. In the case of global warming, even if we knew how large a temperature increase to expect from any particular increase in GHG concentrations, we would not know the resulting cost to society — we cannot predict how a temperature increase would affect agricultural output, land use, etc.² The second type is *ecological uncertainty*, i.e., uncertainty over the evolution of the relevant ecosystems. For example, even if we knew that we could meet a specified policy target for GHG emissions over the next 40 years, we could not predict the resulting levels of atmospheric GHG concentrations at different points in time, nor could we predict the average global equilibrium temperature increase and how that increase would vary regionally.³

Recent studies have begun to examine the implications of irreversibility and uncertainty for environmental policy, at times drawing upon the theory of irreversible investment decisions.⁴ 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, the faster is the rate of learning, the lower is first-period emissions. This is 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 Kolstad

² Likewise, there is considerable uncertainty over the costs of acid rain; even if we could predict the increase in acidity in lakes and rivers from NOX emissions, the impact on fish and other organisms is uncertain, and hence so is the social cost. For most environmental problems there is uncertainty over the future social cost of the environmental degradation, and thus over the social benefit of any policy response.

³ 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, and subject to considerable uncertainty. For a forecasting model of CO₂ emissions with an explicit treatment of forecast uncertainty (see Schmalensee et al., 1998). For general discussions of the uncertainties inherent in the analysis of global warming, see Cline (1992) and Solow (1991). Similar uncertainties exist with respect to acid rain. For example, 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 populations that live in them.

⁴ For 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) and Abel et al. (1996).

goes on to show that if the cost of policy adoption is at least partly sunk, the effect of uncertainty on the initial level of emissions is ambiguous.⁵

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

Finally, related studies by Chao (1995) and Narain and Fisher (1998) deserve mention. Chao examines randomly arriving catastrophic damage from GHG emissions, where the probability of arrival is increasing the emissions rate. Narain and Fisher also develop a model with both ecological and economic irreversibilities. The uncertainty is with respect to the Poisson arrival of a “catastrophe” that drives utility permanently to zero, with a mean arrival rate that is an increasing function of the stock of pollutant. Hence, the probability of a “catastrophe” over any period can be reduced by investing in pollution reduction. This kind of uncertainty has a very limited effect — it simply increases the effective discount rate.

In this paper, I assume that information arrives continually, but there always remains uncertainty over the future evolution of key environmental variables, and over the future costs and benefits of policy adoption. I focus on how irreversibilities and uncertainty interact in affecting the timing and design of policy. The next section begins by laying out the basic analytical framework, and shows how policy design and timing can be treated as an optimal stopping problem. I consider policies, which entail a flow of sunk costs, to reduce emissions of a pollutant

⁵ Kolstad (1996) also obtains this general result in the context of a two-period model. In related work, Hammitt et al. (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. Kolstad (1994) and Kelly and Kolstad (1999) also examined GHG emission policy in the context of a growth model with uncertainty and learning about the value of an unknown parameter. Kolstad finds that temporary emission reduction policies dominate permanent ones, and Kelly and Kolstad characterize the rate of learning in a Bayesian context and show its implications for policy.

⁶ Although he did not do so, Hendricks could also use his model to study the implications of the degree of irreversibility of environmental damage, by varying the parameter that describes the rate of natural GHG removal from the atmosphere. Conrad (1992) also developed a continuous-time model of emission control, in which the social cost of pollution is a quadratic function of the stock of pollutant, with a coefficient that fluctuates as a geometric Brownian motion. The linear–quadratic structure implies that emissions will be zero (a maximum rate) if this coefficient exceeds (is below) a critical value. He shows that this critical value is a declining function of volatility. However, the only irreversibility is with respect to the stock of pollutant, so the results are along the line of those in Henry and Arrow and Fisher.

which accumulates. In Section 3, I consider a model with economic uncertainty, i.e., there is uncertainty over the future social cost of any given stock of pollutant. 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 4, I allow for gradual emission reductions, again in the presence of economic uncertainty. Section 5 examines the implications of ecological uncertainty by allowing the evolution of the stock of pollutant to be stochastic. Section 6 concludes.

2. Analytical framework

In order to get at the basic issues and obtain results that are reasonably easy to interpret, I introduce a model that captures the basic stock externality associated with many environmental problems in as simple a way as possible, while still allowing for 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, 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 , e.g., the rate of CO₂ or SO₂ emissions. I will assume that absent some policy intervention, E_t follows an exogenous trajectory. The evolution of M_t is endogenous, and is given by:

$$dM = [\beta E(t) - \delta M(t)]dt + d\phi(t), \quad (1)$$

where $d\phi(t)$ is the increment of a stochastic process, and δ is the natural rate at which the stock of pollutant dissipates over time. Ignoring the stochastic term this is 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.⁷

I will assume that the flow of social cost associated with the stock variable M_t is specified by a function $B(M_t, \theta_t)$, where θ_t 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 the development of new agricultural techniques that reduce the social cost of a higher M , or alternatively, demographic

⁷ Nordhaus' model supplements Eq. (1) with an adjustment process for temperature:

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

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 and reduced agricultural output. I am simplifying things by dropping the variable T and associating an economic cost directly with M .

changes that raise the cost. One would generally expect $B(M_t, \theta_t)$ to be convex in M_t (at least when M_t is sufficiently large), but for simplicity, I will initially assume that B is linear in M . Uncertainty over the future costs and benefits of policy adoption can be introduced by letting θ follow a stochastic process.

The implications of uncertainty and irreversibility are easiest to see by focusing on policies that are introduced at a specific point in time, and that have a long-term impact on the evolution of E_t , although I will also consider policies that are introduced gradually. Consider a policy introduced at time T that changes the evolution of E_t to some new trajectory E_t^* for $t \geq T$. It would impose a flow of costs on society, some portion of which will be sunk. I 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 of policy characteristics. For example, ω might describe an absolute reduction in E_t , or a reduction in the *expected* rate of growth of E_t .⁸

I will be considering policies that involve 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 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 $K(E_1)$.⁹ The policy objective is to maximize:

$$W = \mathcal{E}_0 \int_0^{\infty} B(M_t, \theta_t) e^{-r't} dt - \mathcal{E}_0 K(E_1) e^{-r\tilde{T}}, \quad (2)$$

subject to Eq. (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 $t = 0$, and r is the discount rate. Thus, we have an optimal stopping problem — we must determine 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 .

This general model with two forms of uncertainty is difficult to solve, so I proceed by considering special cases. In particular, I first allow for only economic uncertainty, and then later I introduce ecological uncertainty.

⁸ For 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 .

⁹ 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.

3. Economic uncertainty

In this section and the next, I allow the social benefit function to shift stochastically over time. I will begin by assuming that policy adoption implies reducing E from its initial level E_0 to zero. I also assume for now that the cost of emissions reduction as a linear function of the size of the reduction, and $B(M_t, \theta_t)$ is a linear function of M_t , i.e.:

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

Hence, if it is optimal to reduce E at all, it is indeed optimal to reduce it to zero. Later I will make the cost of reducing emissions a convex function of the size of the reduction, and let the size of the reduction be a policy choice variable.

Since there is no ecological uncertainty, Eq.(1) is simply:

$$dM/dt = \beta E(t) - \delta M(t).$$

I introduce economic uncertainty by letting θ follow a geometric Brownian motion:

$$d\theta = \alpha\theta dt + \sigma\theta dz. \quad (4)$$

This means the current flow of social cost from a level of pollutant M_t is known, but the future flow of social cost is always uncertain, and the amount of uncertainty grows with the time horizon. Thus, we learn about the social cost of pollution as time passes, but the flow of social cost in the future will always be unknown. Of course one might argue that for some environmental problems, most or all of the uncertainty over social costs will eventually be resolved. (In effect, this means that σ in Eq. (4) will fall over time.) For problems such as global warming, acid rain, and species extinction, there is little evidence of such a resolution of uncertainty (as opposed to a continuing evolution of our assessment of social costs).

Suppose the cost of reducing E from E_0 to zero is given by $K = kE_0$. We want a policy adoption rule that maximizes the net present value function of Eq. (2) subject to Eq. (4) for the evolution of θ , and Eq. (1) for the evolution of M . This problem can be solved using dynamic programming by defining a net present value function for each of two regions. Let $W^N(\theta, M)$ denote the value function for the “no-adopt” region (in which $E_t = E_0$), and let $W^A(\theta, M)$ denote the value function for the “adopt” region (in which $E_t = 0$).

Since $B(M_t, \theta_t) = -\theta_t M_t$, we know $W^N(\theta, M)$ must satisfy the Bellman equation:

$$rW^N = -\theta M + (\beta E_0 - \delta M)W_M^N + \alpha\theta W_\theta^N + \frac{1}{2}\sigma^2\theta^2 W_{\theta\theta}^N. \quad (5)$$

(Partial derivatives are denoted by subscripts, e.g., $W_M^N = \partial W^N / \partial M$). Likewise, $W^A(\theta, M)$ must satisfy the Bellman equation:¹⁰

$$rW^A = -\theta M - \delta MW_M^A + \alpha \theta W_\theta^A + \frac{1}{2} \sigma^2 \theta^2 W_{\theta\theta}^A. \tag{6}$$

These two differential equations must be solved for $W^N(\theta, M)$ and $W^A(\theta, M)$ subject to the following set of boundary conditions:

$$W^N(0, M) = 0, \tag{7}$$

$$W^N(\theta^*, M) = W^A(\theta^*, M) - K, \tag{8}$$

$$W_\theta^N(\theta^*, M) = W_\theta^A(\theta^*, M). \tag{9}$$

Here, θ^* is the critical value of θ at or above which the policy should be adopted. Condition (7) reflects the fact that if θ is ever zero, it will remain at zero thereafter. Condition (8) 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 (9) 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 solution:

$$W^N(\theta, M) = A\theta^\gamma - \frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_0 \theta}{(r - \alpha)(r + \delta - \alpha)}, \tag{10}$$

and:

$$W^A(\theta, M) = -\frac{\theta M}{r + \delta - \alpha}, \tag{11}$$

where A is a positive constant to be determined, and, from boundary condition (7), γ is the positive root of the quadratic equation $1/2 \sigma^2 \gamma(\gamma - 1) + \alpha \gamma - r = 0$, i.e.:

$$\gamma = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left(\frac{\alpha}{\sigma^2} - \frac{1}{2}\right)^2 + \frac{2r}{\sigma^2}} > 1. \tag{12}$$

Note from Eq. (10) that W^N has three components. The first term on the right-hand side of Eq. (10) 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

¹⁰ Eqs. (5) and (6) can be written in more compact form as:

$$rW = -\theta M + (1/dt)\mathcal{L}(dW).$$

Thus the social return on W^N or W^A has two parts, the flow of social cost $-\theta M$, and the expected rate of increase in W (or “capital gain”).

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 results if emissions continued at 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 the flow of cost from emissions E_0 now and in all future periods is $\beta E_0 \theta / (r + \delta - \alpha)(r - \alpha)$.) 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 is 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 (8) and (9):

$$A = \left(\frac{\gamma - 1}{K} \right)^{\gamma - 1} \left[\frac{\beta E_0}{(r - \alpha)(r + \delta - \alpha)\gamma} \right]^\gamma, \tag{13}$$

$$\theta^* = \left(\frac{\gamma K}{\gamma - 1} \right) \left[\frac{(r - \alpha)(r + \delta - \alpha)}{\beta E_0} \right]. \tag{14}$$

Eqs. (13) and (14) apply for any sunk cost of policy adoption, K . If we make use of our assumption that $K = kE_0$, these equations become:

$$A = \left(\frac{\gamma - 1}{k} \right)^{\gamma - 1} \left[\frac{\beta}{(r - \alpha)(r + \delta - \alpha)\gamma} \right]^\gamma E_0, \tag{15}$$

so that the value of the option to adopt the policy is linear in E_0 , and:

$$\theta^* = \left(\frac{\gamma}{\gamma - 1} \right) k(r - \alpha)(r + \delta - \alpha) / \beta. \tag{16}$$

In the absence of any uncertainty or irreversibilities, it would be optimal to adopt the policy when θ reached a level such that:

$$\frac{\beta \theta}{(r - \alpha)(r + \delta - \alpha)} = k. \tag{17}$$

The left-hand side of Eq. (17) is just the present value of the flow of social cost from one extra unit of emissions now and throughout the future (adjusted for the absorption rate β), and the right-hand side is the cost of permanently reducing emissions by one unit. Hence, Eq. (17) is a standard cost–benefit calculation. We can rewrite the equation in terms of a critical value θ^* that triggers policy adoption:

$$\theta^* = k(r - \alpha)(r + \delta - \alpha) / \beta.$$

When there is uncertainty, this critical level θ^* is simply increased by the factor $\gamma / (\gamma - 1)$.

Note that an increase in σ implies a decrease in γ and hence an increase in θ^* . The more uncertainty there is over the future social cost of the pollutant, the greater is the incentive to wait rather than adopt the policy now, and hence the greater must be the current cost in order to trigger adoption. An increase in the discount rate r increases the value of the option to adopt the policy and thus also increases θ^* . The cost, K , is paid in the future when the policy is adopted; hence an increase in r implies a greater reduction in the present value of that cost, so that the option to adopt is worth more but it should be exercised later. An increase in δ , the rate of “depreciation” of the stock of pollutant, also increases θ^* ; a higher value of δ implies that the environmental damage from emissions is more reversible, so that the sunk benefit of adopting the policy now rather than waiting is lower. Finally, θ^* is a monotonically decreasing function of α . The reason is that a higher value of α implies a higher future flow of social cost from whatever the level of M at the time of policy adoption, making it optimal to adopt earlier so that M is smaller.¹¹

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 functions W^N and W^A are 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^* .

We can frame this timing problem in terms of a comparison of the opportunity costs of current adoption with the corresponding opportunity “benefits” by calculating $W^* - W_0$, where W^* is the value function when the adoption decision is made optimally, and W_0 is the value function when the policy is adopted immediately. Suppose $\theta < \theta^*$, so that it is not yet optimal to adopt the policy, and $W^* = W^N$. Since $W_0 = W^A - K$, $W^* - W_0 = W^N - W^A + K$, or:

$$W^* - W_0 = K + A\theta^\gamma - \frac{\beta E_0 \theta}{(r - \alpha)(r + \delta - \alpha)}. \tag{18}$$

The first term on the right-hand side of Eq. (18) 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.

¹¹ Unlike with a financial option, there are no restrictions on the range of α . It would appear from Eq. (16) that $\theta^* \rightarrow 0$ as $\alpha \rightarrow r$, but $\gamma \rightarrow 1$, so $\lim_{\alpha \rightarrow r} \theta^*$ must be found using l’Hopital’s Rule. It is easy to show that θ^* is a continuous function of α at $\alpha = r$.

Note that as the model is currently structured, it would never be optimal to reduce emissions by anything less than 100% (assuming it would be optimal to reduce emissions 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 make K a nonlinear function of the reduction in emissions, and examine policies that involve a one-time partial reduction in emissions, as well as gradual incremental reductions.

3.1. A numerical example

A numerical example will help to explore the characteristics of the solution. Suppose that $\alpha = 0$ (so that the social cost per unit of M is expected to remain constant), $r = 0.04$, $\delta = 0.02$, $\sigma = 0.20$, $\beta = 1$, $E_0 = 300,000$ tons/year, $\theta_0 = \text{US\$}20$ per ton, and $k = 6667$ so that $K = kE_0 = \text{US\$}2$ billion. Then, from Eqs. (12), (13), and (14), $\gamma = 2.0$, $A = 1,953,125$, and $\theta^* = \text{US\$}32$ per ton. Hence, at the current value of $\theta_0 = 20$, the policy should not be adopted. However, the value of the option to adopt it in the future, $A\theta^\gamma$, is $\text{US\$}0.78$ billion. The policy should be adopted when θ reaches $\text{US\$}32$ per ton; at that point $A\theta^\gamma = \text{US\$}2.0$ billion, and boundary conditions (8) and (9) are satisfied. Fig. 1 shows this solution graphically for the case in which $M = 0$ (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 $W^A - K$, and since $M = 0$, $W^A - K$ is a horizontal line at $-K$. (If M were greater than zero, we would have $W^A = -\theta M / (r + \delta - \alpha)$, so we would rotate both the $W^N(\theta)$ curve and the line $W^A - K$ downwards.)




Fig. 1. Solution for $M = 0$.




Fig. 2. θ^* as a function of σ , $\delta = 0.01, 0.02$.

Fig. 2 shows θ^* as a function of σ for $\delta = 0.01$ and 0.02 . Note that 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.

As mentioned in the Introduction, environmental policy debates often focus on the discount rate. However, for many environmental problems, the range of plausible discount rates is much smaller than the range of plausible degrees of uncertainty. It is therefore useful to examine the sensitivity of θ^* to both r and σ . This can be done by calculating the semi-elasticities $d \log \theta^* / dr$ and $d \log \theta^* / d\sigma$, which are given by:

$$\frac{d \log \theta^*}{dr} = \frac{(2r + \delta - 2\alpha)}{(r - \alpha)(r + \delta - \alpha)} - \frac{1}{\gamma(\gamma - 1)\sigma^2 \sqrt{\left(\alpha/\sigma^2 - \frac{1}{2}\right)^2 + 2r/\sigma^2}},$$

and:

$$\frac{d \log \theta^*}{d\sigma} = \frac{(2r + 2\alpha^2/\sigma^2 - \alpha) \left(\left(\alpha/\sigma^2 - \frac{1}{2} \right)^2 + 2r/\sigma^2 \right)^{-1/2} - 2\alpha}{\gamma(\gamma - 1)\sigma^3}.$$

Table 1
Semi-elasticities of θ^*

r	σ	θ^*	$d \log \theta^* / dr$	$d \log \theta^* / d\sigma$
0.04	0.2	32.0	33.33	3.33
0.04	0.4	59.71	27.23	2.89
0.04	0.6	101.47	23.48	2.43
0.04	0.8	158.38	21.25	2.04
0.02	0.2	13.96	52.64	4.47
0.06	0.2	56.56	24.54	2.77
0.08	0.2	87.48	19.47	2.43

Table 1 shows these semi-elasticities, along with θ^* , for different values of r and σ . (In all cases, $\alpha = 0$, $\delta = 0.02$, $K = \text{US}\$2$ billion, and $E_0 = 300,000$ tons/year.) Note that a 0.01 change in r results in approximately the same percentage change in θ^* as does a 0.1 change in σ . But this does not mean that the discount rate is a more important determinant of environmental policy. First, plausible values of the real discount rate are confined to a small range — for analyses of global climate change, for example, between 0.02 and 0.05. But plausible values for σ (or the standard deviation of other stochastic state variables) can fall within a much larger range. Second, most traditional cost–benefit analyses of environmental policy are done by implicitly assuming that $\sigma = 0$. Hence, even if the correct value of σ is only 0.2 or 0.3, just accounting for uncertainty can matter a lot.

3.2. Convex costs and partial reduction in emissions

We now consider policies that only partially reduce emissions. We will assume that the sunk cost of the policy is a quadratic function of the amount that emissions are reduced:

$$K = k_1(E_0 - E_1) + k_2(E_0 - E_1)^2, \quad (19)$$

where $E_0 - E_1$ is the amount of the reduction, and $k_1, k_2 > 0$. Thus, the cost of a 1-unit (permanent) reduction in E is $k(E) = -dK/dE_1 = k_1 + 2k_2(E_0 - E_1)$. We must again 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, i.e., the optimal value of E_1 .

As before, let $W^N(\theta, M)$ and $W^A(\theta, M)$ be the value functions for the “no-adopt” and “adopt” regions, respectively. $W^N(\theta, M)$ must again satisfy the Bellman Eq. (5). However, Eq. (6) is no longer the correct Bellman equation for $W^A(\theta, M)$. After adoption of the policy, $dM/dt = \beta E_1 - \delta M$, so we must include

an additional term, $\beta E_1 W_M^A$, where E_1 is the emissions level after policy adoption. Hence, $W^A(\theta, M)$ now satisfies:

$$rW^A = -\theta M + (\beta E_1 - \delta M)W_M^A + \alpha\theta W_\theta^A + \frac{1}{2}\sigma^2\theta^2 W_{\theta\theta}^A. \tag{20}$$

The solution for $W^A(\theta, M)$ is therefore:

$$W^A(\theta, M) = -\frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_1 \theta}{(r - \alpha)(r + \delta - \alpha)}, \tag{21}$$

while the solution for $W^N(\theta, M)$ is again given by Eq. (10), i.e.:

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

Since E_1 is chosen optimally, it depends on θ at the time of adoption, i.e., on θ^* . Hence, $K = K(E^*(\theta^*))$ in boundary condition (8). Using Eqs. (19) and (21), we choose E^* to maximize the net payoff from policy adoption:

$$\begin{aligned} \max_E [W^A(\theta, M; E) - K(E)] &= -\frac{\theta M}{r + \delta - \alpha} - \frac{\beta E_1 \theta}{(r - \alpha)(r + \delta - \alpha)} \\ &\quad - k_1(E_0 - E) - k_2(E_0 - E)^2, \end{aligned} \tag{22}$$

so that:

$$E^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\theta}{2k_2(r - \alpha)(r + \delta - \alpha)}. \tag{23}$$

We now substitute this expression for E^* into boundary conditions (8) and (9), and then use these conditions to find θ^* and the constant A . Making the substitutions and denoting $\rho \equiv (r - \alpha)(r + \delta - \alpha)$, we find that θ^* must satisfy the quadratic equation:

$$(\gamma - 2)\beta^2\theta^2 - 2\rho(\gamma - 1)\beta k_1\theta + \gamma\rho^2 k_1^2 = 0. \tag{24}$$

We know that γ , which is given by Eq. (12), exceeds 1. However, let us assume for the moment that $\gamma > 2$. Because $W^A(\theta) - W^N(\theta) - K(E^*(\theta))$ is convex in θ , θ^* is the largest root of the quadratic Eq. (24), i.e.:

$$\theta^* = \frac{\rho(\gamma - 1)k_1}{\beta(\gamma - 2)} \left[1 + \sqrt{1 - \frac{\gamma(\gamma - 2)}{(\gamma - 1)^2}} \right]. \tag{25}$$

Then, A is given by:

$$A = \frac{\beta^2}{4k_2\rho^2(\theta^*)^{\gamma-2}} - \frac{\beta k_1}{2k_2\rho(\theta^*)^{\gamma-1}} + \frac{k_1^2}{4k_2(\theta^*)^\gamma}. \tag{26}$$

Given θ^* , we can find $E^*(\theta^*)$ from Eq. (23). It is easy to confirm that as σ increases (so that γ decreases), θ^* will increase and E^* will fall. However, we must account for the fact that E^* must lie between 0 and E_0 .

Let θ_{\max} denote the value of θ for which $E^* = 0$. From Eq. (23), we see that $\theta_{\max} = \rho k_1 / \beta + 2 \rho k_2 E_0 / \beta$. Hence, $E^* \geq 0$ implies that $\theta^* \leq \theta_{\max} = \rho k_1 / \beta + 2 \rho k_2 E_0 / \beta$, or equivalently, that $\gamma \geq 2 + k_1 / k_2 E_0$. Thus, our assumption that $\gamma > 2$, which implies that $\sigma^2 < r - 2\alpha$, will always hold. If σ is sufficiently large so that $\gamma < 2 + k_1 / k_2 E_0$, Eqs. (25) and (26) 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 timing problem is again given by Eqs. (13) and (14) (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 Eqs. (23) and (25) that $E^*(\theta^*) < E_0$ for any $\gamma > 2$.¹²

Once again, a numerical example will be helpful. Suppose that $E_0 = 300,000$ tons/year, $k_1 = 5000$ and $k_2 = 0.0055$ (so that the cost of reducing E to zero would be about US\$2 billion), $\sigma = 0.045$, and as in the earlier example, $\alpha = 0$, $r = 0.04$, $\delta = 0.02$, 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 $\sigma = 0.045$, $\gamma = 6.8$, so that $\theta^* = 17$, i.e., $\theta^* < \theta_{\max}$ so that $E^* > 0$. From Eq. (23), we see that $E^* = 110,606$ tons/year.

The amount that emissions are reduced depends on the degree of uncertainty over the future benefits from a reduction, and on other parameters. Fig. 3 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 $\sigma = 0$, the standard NPV rule 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 slightly. (The reason is that in this numerical example, $\alpha = 0$, so that if $\sigma = 0$, θ cannot rise in the future.) As σ is increased, the critical value θ^* also increases, and E^* falls. Note that for $\sigma > 0.063$, $E^* = 0$, so that θ^* is given by Eq. (14) (with $k_1 E_0 + k_2 E_0^2$ substituted for K) rather than Eq. (25).

We can likewise determine the dependence of θ^* and E^* on other parameters from Eqs. (14), (23), and (25). 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 increases E^* .

¹² It might appear from Eq. (23) 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.




Fig. 3. Partial emission reduction — dependence of E^* and θ^* on σ .

3.3. Convex benefit function

We have assumed that the benefit function $B(M, \theta)$ is linear in M , which makes the optimal policy rules independent of M . This was convenient, but for most environmental problems, the damage from a pollutant is like to rise more than proportionally with the stock of the pollutant. Then the optimal policy rule will depend on the stock. To explore this, we again make the cost of an emission reduction linear in the size of the reduction, and assume emissions must be reduced to zero once a policy is adopted, so that $K = kE_0$. (Of course if the size of the emission reduction could be chosen, $E_1 = 0$ might not be optimal.) Now, however, we let the benefit function $B(M, \theta)$ be quadratic in M :

$$B(M, \theta) = -\theta_t M_t^2. \tag{27}$$

The value functions $W^N(\theta, M)$ and $W^A(\theta, M)$ for the “no-adopt” and “adopt” regions will again satisfy the Bellman Eqs. (5) and (6), but with the term $-\theta M$ replaced by $-\theta M^2$ in each equation. Boundary conditions (7) to (9) also apply. These equations have the following solution:

$$W^N(\theta, M) = A\theta^\gamma - \frac{\theta M^2}{r + 2\delta - \alpha} - \frac{2\beta^2 E_0^2 \theta}{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)} - \frac{2\beta E_0 \theta M}{(r + 2\delta - \alpha)(r + \delta - \alpha)}, \tag{28}$$

and:

$$W^A(\theta, M) = -\frac{\theta M^2}{r + 2\delta - \alpha}, \quad (29)$$

where A is a positive constant to be determined, and γ is again given by Eq. (12). Note that the right-hand side of Eq. (29) and the second term on the right-hand side of Eq. (28) 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 Eq. (28) are the present value of the flow of social cost from *future* emissions at the rate E_0 . The first term on the right-hand side of Eq. (28) is the value of the option to reduce emissions to zero.

Using boundary conditions (8) and (9), the constant A and critical value θ^* are:

$$A = E_0 \left(\frac{\gamma - 1}{k} \right)^{\gamma - 1} \left[\frac{2\beta^2 E_0 + 2\beta(r - \alpha)M}{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)\gamma} \right]^\gamma, \quad (30)$$

and:

$$\theta^* = \frac{(r - \alpha)(r + 2\delta - \alpha)(r + \delta - \alpha)k\gamma}{2\beta(\gamma - 1)[\beta E_0 + (r - \alpha)M]}. \quad (31)$$

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 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 is θ^* . As before, a higher cost of emission reduction, k , and a higher decay rate, δ , lead 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 Eq. (12) as before. Hence, 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 before.

4. 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 and introducing a major new environmental policy, it is unlikely that regulations could be revised 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 provides insight into how the frequency with which regulations can be introduced or changed affects 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 both θ and the stock variable M , and decide when and by how much to mandate emissions reductions in response to changes in these variables.

For this problem to be of interest, either the benefit function or the cost function must be convex. I will assume that the benefit function $B(\theta, M)$ is 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 Eq. (19). Thus, the cost of a 1-unit reduction in E is $\Delta K = k_1 + 2k_2(E_0 - E_1)$. Letting $m_1 = k_1 + 2k_2E_0$ and $m_2 = 2k_2$, the cost of an incremental reduction is:

$$\Delta K = m_1 - m_2 E. \quad (32)$$

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:

$$\Delta M_t = -\frac{\beta \Delta E}{\delta} [1 - e^{-\delta t}], \quad (33)$$

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

$$\Delta W_t = \mathcal{E}_t \int_t^\infty \Delta B_\tau e^{-r(\tau-t)} d\tau = \beta \theta_t \Delta E / \rho, \quad (34)$$

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

This is analogous to the incremental investment and capacity choice problem in Pindyck (1988). Suppose $E_t = E$ currently, and let $W(E; \theta, M)$ be the value function given this E , and given 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:

$$r\Delta F = (\beta E - \delta M)\Delta F_M + \alpha\theta\Delta F_\theta + \frac{1}{2}\sigma^2\theta^2\Delta F_{\theta\theta}, \quad (35)$$

subject to the boundary conditions:

$$\Delta F(E; 0, M) = 0, \quad (36)$$

$$\Delta F(E; \theta^*, M) = \Delta W(\theta^*) - \Delta K(E), \quad (37)$$

$$\Delta F_\theta(E; \theta^*, M) = \Delta W_\theta(\theta^*). \quad (38)$$

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

$$\Delta F = a\theta^\gamma, \quad (39)$$

with $\gamma > 1$ again given by Eq. (12). 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 (37) and (38):

$$\theta^*(E) = \frac{\gamma\rho(m_1 - m_2 E)}{(\gamma - 1)\beta}, \quad (40)$$

$$a = \left(\frac{\beta}{\gamma\rho}\right)^\gamma \left(\frac{\gamma - 1}{m_1 - m_2 E}\right)^{\gamma-1}. \quad (41)$$

To interpret Eq. (40), note that $\rho(m_1 - m_2 E)/\beta$ is the amortized sunk cost of an incremental reduction in emissions, normalized by the absorption rate β . Since $B(\theta, M)$ 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 decay rate δ increases. Finally, given $\theta^*(E)$, we can determine the optimal emissions level E^* .

4.1. Monte Carlo simulation

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/year, the cost function parameters are $k_1 = 5000$ and $k_2 = 0.0055$ (so that the cost of reducing E from 300,000 tons/year to zero would be about US\$2 billion), and $r = 0.04$, $\delta = 0.02$, and $\beta = 1$. I set $\alpha = 0.01$, so that even absent uncertainty, emissions will gradually be reduced as θ increases. I varied σ from 0 to 0.15, in increments of 0.005. For each value of σ , I ran 10,000 simulations of the evolution of θ and the corresponding optimal emissions level E^* .

Fig. 4 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/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 .

Fig. 5 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 $\sigma > 0.13$, 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 time until zero




Fig. 4. Mean optimal emissions level at $t = 0$ and 20 years.

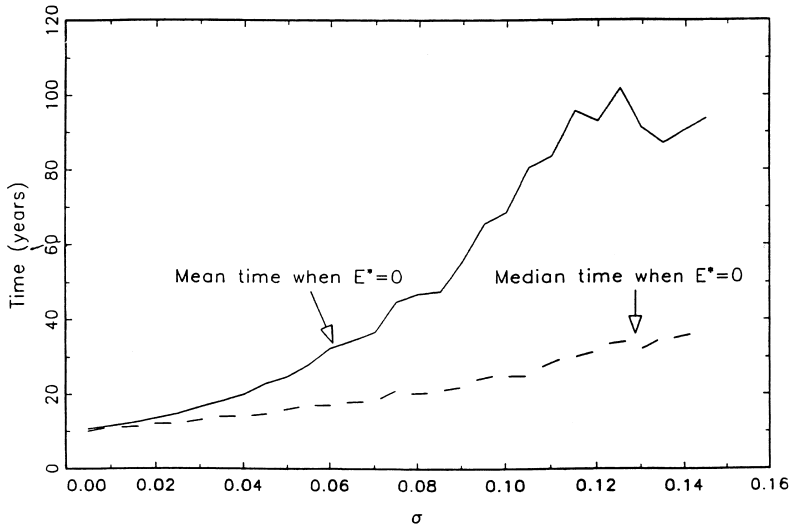


Fig. 5. Mean and median times until emissions are reduced to zero.

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.

5. Ecological uncertainty

So far, the only form of uncertainty that we have considered has been over the parameter θ that shifts the benefit function B . In this section, I will assume that θ remains fixed, but that there is uncertainty over the evolution of M . Specifically, I write Eq. (1) as:

$$dM = (\beta E - \delta M)dt + \sigma dz. \tag{42}$$

Thus, even if the trajectory for E_t were known, future values of M would be uncertain (and normally distributed).¹³

¹³ It might seem more natural to assume that future values of M are lognormally distributed, i.e., to describe the evolution of M by:

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

so that M could never become negative. I use Eq. (42) 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 . The reason is that 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 could not affect the optimal policy. This would be true even if the cost of emission reduction, $K(E)$, were a nonlinear function of the size of the reduction.

We will therefore assume that the benefit function is quadratic in M , i.e., $B(M, \theta) = -\theta M^2$. For simplicity, we will also assume that the (sunk) cost of an emission reduction is linear in the size of the reduction, and that emissions must be reduced to zero once a policy is adopted. Thus, the cost of policy adoption is $K = kE_0$.

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:

$$rW^N = -\theta M^2 + (\beta E_0 - \delta M)W_M^N + \frac{1}{2}\sigma^2 W_{MM}^N, \quad (43)$$

$$rW^A = -\theta M^2 - \delta M W_M^A + \frac{1}{2}\sigma^2 W_{MM}^A. \quad (44)$$

The value functions must also satisfy the boundary conditions:

$$W_M^A(0) = 0, \quad (45)$$

$$W^N(M^*) = W^A(M^*) - K, \quad (46)$$

$$W_M^N(M^*) = W_M^A(M^*), \quad (47)$$

where M^* is the critical value of M that triggers policy adoption. Note that condition (45) applies to the *slope* of $W^A(M)$ at $M=0$, and not the level. Because M follows an arithmetic (rather than geometric) Brownian motion, $M=0$ is not an absorbing barrier, so we do not have a simple boundary condition for the value of $W^A(0)$. However, since $B(M, \theta) = -\theta M^2$, and W^A applies when $E=0$, W^A must reach its maximum at $M=0$.

There is now only one state variable (M), so that Eqs. (43) and (44) are ordinary differential equations. The solution of Eq. (44) and boundary condition (45) for W^A is:

$$W^A(M) = -\frac{\theta M^2}{r + 2\delta} - \frac{\sigma^2 \theta}{r(r + 2\delta)}. \quad (48)$$

This is just the present value of the flow of social cost from the current stock of the pollutant, M , accounting for stochastic fluctuations in M (even when $M=0$). 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 .

Without restrictions on the parameter values, Eq. (43) for W^N does not have an analytical solution. We will see shortly how this equation can be solved numeri-

cally for $W^N(M)$ and, simultaneously, for the critical value M^* . First we will consider the special case of $\delta = 0$, for which $W^N(M)$ does have an analytical solution.

5.1. Environmental damage completely irreversible ($\delta = 0$)

When $\delta = 0$, the solution for $W^A(M)$ in Eq. (48) reduces to:

$$W^A(M) = -\theta(rM^2 + \sigma^2)/r^2. \tag{49}$$

As for Eq. (43), the homogeneous solution has the form $W^N = B_1e^{\phi_1M} + B_2e^{\phi_2M}$, and the particular solution is a quadratic in M . By direct substitution, the solution for W^N is:

$$W^N(M) = Be^{\phi M} - \frac{\theta(rM^2 + \sigma^2)}{r^2} - \frac{2\beta E_0\theta(\beta E_0 + rM)}{r^3}, \tag{50}$$

where:

$$\phi = -\frac{\beta E_0}{\sigma^2} \left[1 - \sqrt{1 + 2r\sigma^2/\beta^2 E_0^2} \right] > 0, \tag{51}$$

and B is a constant that must be determined. The first term on the right-hand side of Eq. (50) is the value of the option to adopt the policy. The second term is the present value of the flow of social cost from the current stock of the pollutant, M , allowing for stochastic fluctuations in M in the future. 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, again accounting for stochastic fluctuations in M (which now also has a deterministic component of growth).

The constant B and the critical value M^* can be found from the solution for $W^A(M)$ along with the value-matching and smooth-pasting conditions (46) and (47):

$$B = \frac{2\beta E_0\theta}{r^2\phi} e^{-\phi M^*} > 0, \tag{52}$$

and:

$$M^* = -\frac{\beta E_0}{r} - \frac{\sigma^2}{\beta E_0(1 - \sqrt{1 + 2r\sigma^2/\beta^2 E_0^2})} + \frac{r^2 K}{2\beta E_0\theta}. \tag{53}$$

This solution has properties that we would expect. Note in particular that $\partial M^*/\partial K > 0$, $\partial M^*/\partial r > 0$, $\partial M^*/\partial \theta < 0$, and $\partial M^*/\partial \sigma > 0$. Thus, stochastic fluctuations in M create an incentive to delay policy adoption. As a numerical example, and for comparison to results shown below for the more general case of $\delta > 0$, we will set $r = 0.04$, $K = 4$, $E_0 = 0.3$, $\beta = 1$, and $\theta = 0.002$. Then, if $\sigma = 0$, the policy should be adopted immediately for any (positive) value of M .




Fig. 6. Solution method for stochastic M .

However, if $\sigma = 1$, the policy should only be adopted when $M \geq M^* = 6.74$, and when $\sigma = 4$, the policy should be adopted when $M \geq M^* = 16.21$.

5.2. General case

For the more general case of $\delta > 0$, $W^N(M)$ and the critical value M^* must be found numerically. This is done by utilizing the solution for $W^A(M)$ given by Eq. (48), along with the boundary conditions (46) and (47). To obtain a numerical solution, we begin with a candidate number (e.g., a best guess) for M^* ; call this M_0^* . We then use Eqs. (46–48) to get $W^N(M_0^*)$ and $W_M^N(M_0^*)$, and we solve Eq. (43) backwards to determine a corresponding candidate solution for $W^N(M)$ for all M between 0 and M_0^* . To be the actual solution, the candidate solution must satisfy one regularity condition for all values of M between 0 and M_0^* , and a second condition at $M = 0$. First, since an increase in M always implies a reduction in the flow of current and future social benefits, we must have $W_M^N < 0$ for all values of M between 0 and M^* . Second, because of the convexity of $B(M, \theta)$, we must have $W_{MM}^N < 0$ at $M = 0$.¹⁴ Thus, the candidate number for M^* is repeatedly adjusted up or down (in smaller and smaller steps) until both of these conditions are satisfied.

¹⁴ If $\delta = 0$, we will have $W_{MM}^N \leq 0$ for all values of M between 0 and M^* . However, if $\delta > 0$, this condition must apply only at $M = 0$, since δM is increasing in M .

Table 2

Solutions for M^* ($r = 0.04$, $\beta = 1$, $K = 4$, $E_0 = 0.3$, $\theta = 0.002$)

σ	$\delta = 0$	$\delta = 0.02$
0.3	5.48	11.08
0.5	5.73	11.59
0.8	6.28	12.45
1.0	6.74	13.05
2.0	9.59	16.47
4.0	16.21	25.75

5.3. Numerical example

This solution method is easiest to see in the context of a numerical 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 K in billions of dollars. Since the benefit function is $B = -\theta M^2$, we measure θ in billion dollars/(million tons)². We set $K = 4$, $E_0 = 0.3$, $\theta = 0.002$, $\sigma = 1$, $\alpha = 0$, and, as before, $r = 0.04$, $\delta = 0.02$, and $\beta = 1$.

As Fig. 6 shows the solution for M^* in this example is 13.05. The figure shows candidate solutions for $W^N(M)$ corresponding to different values of M^* , along 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 of M^* above 13.05, $W_{MM}^N(M) > 0$ for small M . The solution procedure searches over candidate values of M^* , using an increasingly narrow range.

Table 2 shows the critical value M^* for values of σ ranging from 0.3 to 4.0, and for δ equal to 0 and 0.02. The table shows that M^* increases with σ , but it also shows how M^* increases with δ . A higher δ implies a lower (or negative) rate of drift for M — emissions are more reversible, so the present value of the flow of social cost for any current value of M is lower, and a higher M is needed to justify the sunk cost of policy adoption.

6. Conclusions

I have focused largely on a one-time policy adoption to reduce emissions of a pollutant. If the policy imposes sunk costs on society, and if it can be delayed, there is an opportunity cost of adopting the policy now rather than waiting for more information. This is analogous to the incentive to wait that arises with irreversible investment decisions. In the case of environmental policy, however, this opportunity cost must be balanced against the opportunity “benefit” of early action — a reduced stock of pollutant that might decay only slowly, imposing irreversible or nearly irreversible costs on society.

In the simple models presented in this paper, an increase in uncertainty, whether over future costs and benefits of reduced emissions, or over the evolution of the stock of pollutant, leads to a higher threshold for policy adoption. This is because policy adoption involves a sunk cost associated with a discrete reduction in the entire trajectory of future emissions, whereas inaction over any small time interval only involves continued emissions over that interval. This is true even for gradual emission reductions — a small reduction is a reduction in the entire trajectory. Hence, in my framework greater uncertainty always leads to greater delay, although the effect is smaller the smaller is the decay rate, δ .

The validity of this result depends on the extent to which environmental policy is indeed irreversible, in the sense of involving commitments to future flows of sunk costs. It seems to me that this kind of irreversibility is often an inherent aspect of environmental policy, both for policies that are in place (e.g., the Clean Air Act), and for policies under debate (e.g., GHG emission reductions). Nonetheless, the assumption of complete irreversibility made in this paper (i.e., all costs of policy adoption are sunk) may be extreme. Richer models are needed to explore the implications of relaxing this assumption somewhat.

In the models presented here, economic and ecological uncertainty were treated separately. Ideally, we would like to allow both θ and M to evolve stochastically at the same time. This can be done, but then the value functions $W^A(\theta, M)$ and $W^N(\theta, M)$ will satisfy more complicated partial differential equations that must be solved numerically. Solution methods for such models are discussed in Pindyck (1998), but are beyond the scope of this paper.

Acknowledgements

This research was supported by the National Science Foundation under Grant No. SES90-22823, by M.I.T.'s Center for Energy and Environmental Policy Research, and by the M.I.T. Program on Global Climate Change. My thanks to Avinash Dixit, Scott Farrow, Charles Kolstad, William Nordhaus, Ian Parry, Stephen Peck, V. Kerry Smith, Martin Weitzman, and an anonymous referee for helpful comments and suggestions, and to Eileen Brooks, Tomomi Kumagai, Stephen McNamara, and Sung-Jun Woo for their superb research assistance.

References

- Abel, A.B., Dixit, A.K., Eberly, J.C., Pindyck, R.S., 1996. Options, the value of capital, and investment. *Quarterly Journal of Economics*, 753–777.
- Arrow, K.J., Fisher, A.C., 1974. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics* 88, 312–319.
- Chao, H.-P., 1995. Managing the risk of global climate catastrophe: an uncertainty analysis. *Risk Analysis* 15, 69–78.

- Cline, W.R., 1992. *The Economics Of Global Warming*. Institute for International Economics, Washington, D.C.
- Conrad, J.M., 1992. Stopping rules and the control of stock pollutants. *Natural Resource Modeling* 6, 315–327, (summer).
- Dixit, A., 1992. Investment and hysteresis. *Journal of Economic Perspectives* 6, 107–132, (winter).
- Dixit, A., Pindyck, R.S., 1994. *Investment Under Uncertainty*. Princeton Univ. Press, Princeton, NJ.
- Fisher, A.C., Hanemann, W.M., 1990. Information and the dynamics of environmental protection: the concept of the critical period. *Scandinavian Journal of Economics* 92, 399–414.
- Hammit, J.K., Lempert, R.J., Schlesinger, M.E., 1992. A sequential-decision strategy for abating climate change. *Nature* 357, 315–318, (May 28).
- Hendricks, D., 1992. *Optimal Policy Response to an Uncertain Threat: The Case of Global Warming*, unpublished manuscript, Kennedy School of Government, Harvard University, March.
- Henry, C., 1974. Investment decisions under uncertainty: the irreversibility effect. *American Economic Review* 64, 1006–1012, (December).
- Kelly, D.L., Kolstad, C.D., 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23, 491–518.
- Kolstad, C.D., 1992. *Regulating a Stock Externality Under Uncertainty with Learning*, Working Paper No. 92-0112, Department of Economics, University of Illinois at Urbana-Champaign, March.
- Kolstad, C.D., 1994. *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, January.
- Kolstad, C.D., 1996. Fundamental irreversibilities in stock externalities. *Journal of Public Economics* 60, 221–233.
- Krutilla, J.V., Fisher, A.C., 1975. *The Economics of Natural Environments*. Johns Hopkins Univ. Press, Baltimore.
- Narain, U., Fisher, A.C., 1998. *Irreversibility, Uncertainty, and Catastrophic Global Warming*, Working Paper No. 843, Department of Agricultural and Resource Economics, University of California at Berkeley, June.
- Nordhaus, W.D., 1991. To slow or not to slow: the economics of the greenhouse effect. *The Economic Journal* 101, 920–937, (July).
- Pindyck, R.S., 1988. Irreversible investment, capacity choice, and the value of the firm. *American Economic Review* 79, 969–985, (December).
- Pindyck, R.S., 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature* 29, 1110–1152, (September).
- Pindyck, R.S., 1998. *Sunk Costs and Sunk Benefits in Environmental Policy*, Working Paper No. 95-003WP, M.I.T. Center for Energy and Environmental Policy, December.
- Schmalensee, R., Stoker, T.M., Judson, R.A., 1998. World energy consumption and carbon dioxide emissions: 1950–2050. *Review of Economics and Statistics* 53, 15–27, (February).
- Solow, A.R., 1991. Is there a global warming problem? In: Dornbusch, R., Poterba, J. (Eds.), *Global Warming: Economic Policy Responses*. MIT Press, Cambridge, MA.