



Optimal timing problems in environmental economics

Robert S. Pindyck*

Massachusetts Institute of Technology, Sloan School of Management, E52-453, Cambridge, MA 02139, USA

Abstract

Because of the uncertainties and irreversibilities that are often inherent in environmental degradation, its prevention, and its economic consequences, environmental policy design can involve important problems of timing. I use a simple two-period model to illustrate these optimal timing problems and their implications for environmental policy. I then lay out and solve a continuous-time model of policy adoption in which the policy itself entails sunk costs, and environmental damage is irreversible. The model generalizes earlier work in that it includes two stochastic state variables; one captures uncertainty over environmental change, and the other captures uncertainty over the social costs of environmental damage. Solutions of the model are used to show the implications of these two types of uncertainty for the timing of policy adoption. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: Q28; L51; H23

Keywords: Environmental policy; Irreversibilities; Optimal stopping; Uncertainty; Option value; Stock pollutants; Global warming

1. Introduction

Optimal timing (or “stopping”) problems are an important class of stochastic control problems that arise in economics and finance, as well as other fields. Unlike “continuous control” problems, in which one or more control variables are adjusted continuously and optimally over time to maximize some objective function, these problems involve

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

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

the optimal timing of a discrete action.¹ Important examples include optimal exercise rules for financial options (e.g., finding the threshold price of a dividend-paying stock at which it is optimal to exercise a call option on that stock), and optimal capital investment and disinvestment decisions (e.g., finding the threshold prices of copper at which it is optimal to shut down an existing copper mine or invest in a new mine).²

As illustrated by a small but growing literature, optimal timing problems of this sort also arise in environmental economics. These problems are of the following basic form: At what point should society adopt a (costly) policy to reduce emissions of some environmental pollutant? The traditional approach to this problem applies standard cost–benefit analysis (a simple NPV rule in capital budgeting terms), and would thus recommend adopting a policy if the present value of the expected flow of benefits exceeds the present value of the expected flow of costs. This standard approach, however, ignores three important characteristics of most environmental problems. 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 emissions of greenhouse gases (GHG) such as CO₂, nor do we know the economic impact of higher temperatures. Second, there are usually important irreversibilities associated with environmental policy. These irreversibilities can arise not only 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. These uncertainties, irreversibilities, and the possibility of delay can significantly affect the optimal timing of policy adoption.

There are two kinds of irreversibilities, and they work in opposite directions. First, an environmental policy imposes *sunk costs* on society. For example, coal-burning utilities might be forced to install scrubbers or pay more for low-sulfur coal, or firms might have to scrap existing machines and invest in more fuel-efficient ones. In addition, political constraints may make an environmental policy itself difficult to reverse, so that these sunk costs are incurred over a long period of time, even if the original rationale for the policy disappears. These kinds of sunk costs create an opportunity cost of adopting a policy now, rather than waiting for more information, and this biases traditional cost–benefit analysis in favor of policy adoption. Second, environmental damage can be partially or totally irreversible. For example, increases in GHG concentrations are long lasting, and the damage to ecosystems from higher global temperatures (or from acidified lakes and streams, or the clear-cutting of forests) can be permanent. Thus, 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.³

¹ Kendrick (1981) provides a textbook treatment of what I have termed “continuous control” problems. He gives particular attention to stochastic adaptive control problems (in which optimal feedback rules are found for the response of control variables to stochastic shocks in the state variables), as well as “dual control” problems (in which control variables are adjusted to obtain information as well as directly the trajectories of the state variables).

² For a textbook treatment of such optimal capital investment decisions, see Dixit and Pindyck (1994).

³ This point was made some two decades ago by Arrow and Fisher (1974), Henry (1974), and Krutilla and Fisher (1975).

There are also two types of uncertainties that 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, we would not know the resulting cost to society—we cannot predict how a temperature increase would affect agricultural output, land use, etc. The second 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 would not know the resulting levels of atmospheric GHG concentrations and average global equilibrium temperature increase.⁴

A number of 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. I will not attempt to survey this literature here.⁵ Instead, I will examine the optimal timing of environmental policy in two ways.

First, I lay out a simple two-period model, in which the choice is whether to adopt an emissions-reducing policy now, or wait some fixed period of time (e.g., 20 years), and then, depending on new information that has arrived regarding the extent of environmental degradation and its economic cost, either adopt the policy or reject it. Although this model is very restrictive and in some ways unrealistic, it brings out many of the key insights.⁶

Second, I extend and generalize the continuous-time model of environmental policy adoption in Pindyck (2000). In that model, an emissions-reducing policy can be adopted at any time. Information arrives continually, but there is always uncertainty over the future evolution of key environmental variables, and over the future costs and benefits of policy adoption. As in this paper, I focused on how irreversibilities and uncertainty interact in affecting the timing of policy adoption. However, in that earlier work, I included only one form of uncertainty at a time—economic or ecological—but not both together. Here, I generalize the model to include both forms of uncertainty at the same time. This provides additional insight into their individual effects on policy adoption, as well as the effects of their interactions. In particular, I show that once the stock of pollutant becomes moderately large, uncertainty over its future growth matters much less than economic uncertainty for optimal policy adoption.

In the next section, I lay out the basic two-period model of policy adoption. Although it is quite simple, the model illustrates how and why uncertainty affects the timing and design of an emissions-reducing policy. In Section 3, I present the continuous-time

⁴For 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 NO_x emissions will affect the future acidity of lakes and rivers, or the viability of the fish populations that live in them.

⁵Examples of this literature include Conrad (1992), Hendricks (1992), Kelly and Kolstad (1999), Kolstad (1996), Narain and Fisher (1998), and Pindyck (1996, 2000).

⁶Hammitt 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. Another related study is Peck and Teisburg (1992).

model and show how it can be solved. By calculating solutions for different combinations of parameter values, I show how economic and ecological uncertainties affect the optimal timing of policy adoption. Section 4 concludes.

2. A two-period model

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) it should be adopted. In other words, adopting a policy today competes not only with never adopting the policy, but also with adopting it next year, in two years, and so on. Thus, the policy problem is one of optimal stopping.

As in Pindyck (1996, 2000), I will work with a bare-bones 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, e.g., the average concentration of CO₂ in the atmosphere or the acidity level of a lake. Let E_t be a flow variable that controls M_t . For example, E_t might be the rate of CO₂ or SO₂ emissions. We will assume that in the absence of some policy intervention, E_t follows an exogenous 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 dissipates over time.

I will assume that the flow of social cost (i.e., negative benefit) associated with the stock variable M_t can be specified by a function $B(M_t, \theta_t)$, where θ_t shifts stochastically over time reflecting changes in tastes and technologies. For example, if M is the GHG concentration, shifts in θ might reflect the arrival of new agricultural techniques that reduce the social cost of a higher M , or demographic changes that raise the cost. One would generally expect $B(M_t, \theta_t)$ to be convex in M_t , but for simplicity I will assume in this section that B is linear in M :

$$B(M_t, \theta_t) = -\theta_t M_t. \quad (2)$$

I also begin with a restrictive assumption about the evolution of E_t : Until a policy is adopted, E_t stays at the constant initial level E_0 , and policy adoption implies a once-and-for-all reduction to a new and permanent level E_1 , with $0 \leq E_1 \leq E_0$. Finally, I assume that the social cost of adopting this policy is completely sunk, and its present value at the time of adoption, which I denote by $K(E_1)$, is a function of the size of the emission reduction.

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

$$W = \mathcal{E}_0 \int_0^{\infty} B(M_t, \theta_t) e^{-rt} dt - \mathcal{E}_0 K(E_1) e^{-r\bar{T}} \quad (3)$$

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.

In this section, I make T exogenous. Thus the choices are to adopt the policy today (making M_T smaller than it would be otherwise), or to wait until time T and then, after evaluating the situation, decide whether or not to adopt the policy. I will also assume initially that if the policy is adopted, emissions are reduced from E_0 to zero. Hence, the sunk cost of policy adoption is simply a number, K . (Later in this section I will consider the possibility of reducing E to some level $E_1 > 0$, and I will also examine the adoption decision when the policy is partially reversible.)

For this problem to be interesting, we need to introduce some source of uncertainty. I will assume that there is economic 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, given that we have waited until time T to make a decision.)

By solving Eq. (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,

$$M_t = \begin{cases} (\beta E_0 / \delta)(1 - e^{-\delta t}) + M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T, \\ (\beta E_0 / \delta)(e^{\delta T} - 1)e^{-\delta t} + M_0 e^{-\delta t} & \text{for } t > T, \end{cases} \tag{4}$$

where M_0 is the initial value of M_t . If the policy is never adopted, the first line of Eq. (4) 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 :

$$W_N = - \int_0^\infty \theta_0 M_t e^{-rt} dt = - \frac{\theta_0 M_0}{(r + \delta)} - \frac{\beta E_0 \theta_0}{r(r + \delta)}. \tag{5}$$

Next, suppose the policy is adopted at time $t = 0$. Then a sunk cost K is incurred immediately, $E_t = 0$ always, and the value function is

$$W_0 = - \frac{\theta_0 M_0}{r + \delta} - K. \tag{6}$$

A conventional cost–benefit analysis would recommend adoption of the policy if the net present value $W_0 - W_N$ is positive, i.e., if $\beta E_0 \theta_0 / r(r + \delta) - K > 0$.

Let us introduce some numbers so that we can compare these two alternatives: the present value of the cost to society of policy adoption, K , is \$2 billion, $r = 0.04$, $\delta = 0.02$, $\beta = 1$ (i.e., all emissions are completely absorbed into the ecosystem),

Table 1
Parameter values

Parameter	Value
r (discount rate)	0.04
δ (pollutant decay rate)	0.02
β (absorption factor)	1
K (PV of cost of policy adoption)	\$2 billion
E_0 (emission rate)	300,000 tons/yr
θ_0 (current social cost)	\$20/ton/yr
$\underline{\theta}$ (future social cost, low)	\$10/ton/yr
$\bar{\theta}$ (future social cost, high)	\$30/ton/yr
T (fixed delay time)	10 yr

$E_0 = 300,000$ tons/yr, and $\theta_0 = \$20/\text{ton}/\text{yr}$.⁷ In what follows, I will also assume that $\underline{\theta} = \$10/\text{ton}/\text{yr}$, and $\bar{\theta} = \$30/\text{ton}/\text{yr}$. These parameter values are summarized in Table 1.

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 and then adopt the policy only if $\theta_T = \bar{\theta}$. Denoting the value function that corresponds to this course of action by W_T , using Eq. (4), and noting that the probability that $\theta_T = \bar{\theta}$ is 0.5, we have

$$W_T = -\frac{\theta_0}{r + \delta} \left(M_0 + \frac{\beta E_0}{r} \right) + \frac{\beta E_0}{r(r + \delta)} \left(\theta_0 - \frac{1}{2} \underline{\theta} \right) e^{-rT} - \frac{1}{2} K e^{-rT}. \quad (7)$$

Is it better to adopt the policy at time $t = 0$ or wait until T ? Comparing W_0 to W_T :

$$\Delta W_T = W_T - W_0 = K \left(1 - \frac{1}{2} e^{-rT} \right) - \frac{\beta E_0 \theta_0}{r(r + \delta)} (1 - e^{-rT}) - \frac{\beta E_0 \underline{\theta}}{2r(r + \delta)} e^{-rT}. \quad (8)$$

It is better to wait until time T if and only if $\Delta W_T > 0$.

This expression for ΔW_T has three components. The first term on the right-hand side of Eq. (8) is the present value of the net expected cost savings from delay; the sunk cost K is initially avoided, and there is only a 0.5 probability that it will have to be incurred at time T . Hence, this term represents the opportunity cost of adopting the policy now rather than waiting. The second and third terms are the present value of the expected increase in social cost from environmental damage due to delay. The second term is the cost of additional pollution between now and time T that results from delay, and the last term—the probability that $\theta_T = \underline{\theta}$, times the present value of the cost of 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 represent an “opportunity benefit” of adopting the policy now.

⁷ I am implicitly assuming that the discount rate r is the real risk-free rate of interest, so a value of 0.04 is reasonable. A value of 0.02 for δ is high for the rate of natural removal of atmospheric GHGs (a consensus estimate would be closer to 0.005), but is low for acid concentrations in lakes.

We can therefore rewrite Eq. (8) as

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

where

$$F_C = K(1 - \frac{1}{2}e^{-rT}) \quad (9)$$

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

$$F_B = \frac{\beta E_0 \theta_0}{r(r + \delta)}(1 - e^{-rT}) + \frac{\beta E_0 \theta}{2r(r + \delta)} e^{-rT} \quad (10)$$

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. (As $\delta \rightarrow \infty$, environmental damage becomes completely reversible, and $F_B \rightarrow 0$.) An increase in the discount rate, r , increases F_C and reduces 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 T is 10 yr. Substituting this and the other base case parameter values into Eqs. (9) and (10) 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. In this case the opportunity cost of current adoption slightly outweighs the opportunity benefit.

We assumed that if we delayed the adoption decision until time T , it would then 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 T , this is just the value of θ for which $W_0 - W_N$ is zero. Using Eqs. (6) and (5), we see that this value is given by

$$\hat{\theta}_T = r(r + \delta)K/\beta E_0. \quad (11)$$

For our base case parameter values, $\hat{\theta}_T = \$16/\text{ton}/\text{yr}$. 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 what the optimal amount of reduction should be. However, $B(M_t, \theta_t)$ is linear in M_t and M_t depends linearly on E (see Eq. (4)), so 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 is independent of the size of the reduction. This will not be the case 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, as discussed below.

2.1. Irreversibility, uncertainty, and a “good news principle”

We assumed that the cost of policy adoption 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 pollutant decay rate is smaller than assumed earlier—specifically, that δ is 0.01 instead of 0.02. 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 opportunity benefit of current adoption greater than the opportunity 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 opportunity cost of adopting now, because there is still a 0.5 probability that at time T we will regret having made the decision to spend K and adopt the policy; F_C is \$1.330 billion as before. However, this increase in variance reduces the opportunity 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 0.02, we have $F_B = 0.824 + 0 = \$0.824$ billion, so that $\Delta W_T = 1.330 - 0.824 = \0.506 billion, which is much larger than before. Even if we lower δ to 0.01 (so that environmental damage is more irreversible), $F_B = 0.989$, $\Delta W_T = \$0.341$ billion, and it is still optimal to wait.

This 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., that $\theta_T = \bar{\theta}$, 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. But because 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.

2.2. Allowing for policy reversal

So far we have assumed that once a policy to reduce emissions to zero has been adopted, it would remain in place indefinitely. We now examine how the timing decision changes when a policy adopted at time 0 can be at least partially reversed at time T . In effect, we will be relaxing our earlier assumption that the cost of policy adoption is completely sunk.

We will assume that upon reversal, a fraction ϕ of the cost K can be recovered. This 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. (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 an awareness.)

We again assume that θ_T will equal $\underline{\theta}$ or $\bar{\theta}$, each with probability 0.5. 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 reverse it at time T if $\theta_T = \underline{\theta}$? Clearly, this will depend on the value of ϕ , i.e., the fraction of K that can be recovered.

As before, let W_0 denote the value function when we adopt the policy at time 0, but note that it is now different because of the possibility of policy reversal. Specifically, W_0 must now include the value of society's option (a *put* option) to reverse the policy at time T and recover ϕK . Also, let W_T again 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 policy reversal *after* time T .)

To determine W_0 in this case, we need the trajectory for M_t when the policy is adopted at $t = 0$ and reversed at $t = T$. From Eq. (1), that trajectory is given by

$$M_t = \begin{cases} M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T, \\ (\beta E_0 \underline{\theta} / \delta) [1 - e^{-\delta(t-T)}] + M_0 e^{-\delta t} & \text{for } t > T. \end{cases} \quad (12)$$

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

$$(\beta E_0 \underline{\theta} / \delta) \int_T^\infty [1 - e^{-\delta(t-T)}] e^{-r(t-T)} dt < \phi K. \quad (13)$$

This implies that the policy should be reversed if $\theta_T = \underline{\theta}$ at time T as long as

$$\phi > \phi_{\min} = \frac{\beta E_0 \underline{\theta}}{r(r + \delta)K}. \quad (14)$$

For our numerical example, with $E_0 = 300,000$ tons/yr, $K = \$2$ billion, and $\underline{\theta} = \$10$ /ton/yr, $\phi_{\min} = 0.625$. Thus if $\phi < 0.625$, the option to reverse the policy at time T has no value, and our earlier results still hold.

Suppose $\phi > \phi_{\min}$, so that the policy would indeed be reversed if $\theta_T = \underline{\theta}$. Although W_T is still given by Eq. (7), by using Eq. (12) we can see that W_0 is now given by

$$W_0 = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_0 \underline{\theta}}{2r(r + \delta)} e^{-rT} + \frac{1}{2} \phi K e^{-rT} - K. \quad (15)$$

The second and third terms on the right-hand side of (15) represent the value of the option to reverse the policy at time T . That option value is positive as long as $\phi > \phi_{\min}$.

Using Eqs. (7) and (15), we find that $\Delta W_T = W_T - W_0$ is now given by

$$\Delta W_T = K [1 - \frac{1}{2}(1 + \phi)e^{-rT}] - \frac{\beta E_0 \theta_0}{r(r + \delta)} (1 - e^{-rT}). \quad (16)$$

The first term on the right-hand side of (16) is the opportunity cost of early policy adoption, which we have denoted by F_C , and the second term is the opportunity benefit, F_B . Comparing Eqs. (16) and (8), note that both F_C and F_B are now smaller. Compared to the case where the policy cannot be reversed, F_C is reduced by the amount $\frac{1}{2}\phi K e^{-rT}$, which is the expected value of the portion of sunk cost that can be recovered. In addition, F_B no longer has the term in $\underline{\theta}$, because now if $\theta_T = \underline{\theta}$, the policy will be reversed.

Returning to our numerical example, suppose that $\phi = 0.9$, which exceeds ϕ_{\min} . Then,

$$\Delta W_T = F_C - F_B = \$0.726 \text{ billion} - \$0.824 \text{ billion} = -\$0.098 \text{ billion},$$

so that immediate adoption is better than waiting. The reason is that while the option to reverse the policy has reduced both F_C and F_B , it has reduced F_C by more. (F_C falls from \$1.33 billion to \$0.73 billion, a change of \$0.60 billion, and F_B falls from \$1.24 billion to \$0.82 billion, a change of \$0.42 billion.)

Suppose we increase the variance of θ_T as we did before by letting $\underline{\theta}$ and $\bar{\theta}$ equal 0 and 40, respectively, rather than 10 and 30. If $\phi = 0.9$, $\Delta W_T = -\$97.8$ 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 reversal is optimal if $\theta_T = \underline{\theta}$. From Eq. (14), we see that now $\phi_{\min} = 0$, so that once the policy has been adopted, reversal is always optimal if $\theta_T = \underline{\theta}$. 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 $\phi = 0.1$, $\Delta W_T = \$438.4$ 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 Eq. (16), that value is $\phi = 0.754$. For $\phi > 0.754$, the put option is sufficiently valuable so that early adoption is economical.

Although $\underline{\theta}$ does not appear in Eq. (16), 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 only $\underline{\theta}$ affects whether we would indeed exercise the put option should this low value of θ_T be realized. This is another example of the “good news principle” discussed earlier.

2.3. 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 further by allowing for a partial reduction in emissions. This is of interest only if the cost of policy adoption is a convex function of the amount of emission reduction (or, alternatively, if the benefit function $B(M_t, \theta_t)$ is convex in M_t). Suppose that the cost of (permanently) reducing E from E_0 to $E_1 \geq 0$ is

$$K = k_1(E_0 - E_1) + k_2(E_0 - E_1)^2 \quad (17)$$

with $k_1, k_2 \geq 0$. Then the marginal cost of reducing E an additional unit below E_1 is

$$k(E) = -\frac{dK}{dE} = k_1 + 2k_2(E_0 - E_1). \quad (18)$$

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 Eq. (1) to determine the trajectory for 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 that $E_t = E_1 \geq 0$ for $t \geq T$, so the trajectory for M_t is given by⁸

$$M_t = \begin{cases} (\beta E_0 / \delta)(1 - e^{-\delta t}) + M_0 e^{-\delta t} & \text{for } 0 \leq t \leq T, \\ (\beta E_0 / \delta)(e^{\delta T} - 1)e^{-\delta t} + (\beta E_1 / \delta)[1 - e^{-\delta(t-T)}] + M_0 e^{-\delta t} & \text{for } t > T. \end{cases} \quad (19)$$

First, suppose we reduce E from E_0 to an arbitrary level E_1 at $t = 0$. Then the value function, which we will denote by $W_0(E_1)$, is

$$W_0(E_1) = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_1 \theta_0}{r(r + \delta)} - K(E_1). \quad (20)$$

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

$$W_0(E_1) - W_N = \frac{\beta(E_0 - E_1)\theta_0}{r(r + \delta)} - K(E_1). \quad (21)$$

If we indeed adopt the policy at $t = 0$, we will choose E_1 to maximize this NPV. Using Eq. (17) for $K(E_1)$, the optimal value of E_1 is

$$E_1^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\theta_0}{2k_2 r(r + \delta)}. \quad (22)$$

Setting $E_1 = E_1^*$, the NPV of immediate adoption becomes

$$W_0(E_1^*) - W_N = \frac{1}{4k_2} \left[\frac{\beta\theta_0}{r(r + \delta)} - k_1 \right]^2. \quad (23)$$

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

A numerical example is again helpful. We will use the same parameter values as before (see Table 1), and set $k_1 = 4000$ and $k_2 = 0.02$ (so that reducing E from 300,000 tons/yr to zero would cost \$3.0 billion). In this case, $E_1^* = 191,667$ tons/yr, so that $\Delta E^* = E_0 - E_1^* = 108,333$ tons/yr, $K(\Delta E^*) = \$0.668$ billion, and the NPV of immediate policy adoption is $W_0(E_1^*) - W_N = \$0.234$ billion.

So far we have compared reducing emissions to some amount E_1 at time 0 to never reducing them. Suppose instead that we wait until time T 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 reduce emissions less, to $\underline{E} > \bar{E}$. Using Eq. (19) for M_t and for the time being letting E and \bar{E} be arbitrary, we can determine that the value function $W_T(\underline{E}, \bar{E})$ is

$$W_T(\underline{E}, \bar{E}) = -\frac{\theta_0 M_0}{r + \delta} - \frac{\beta E_0 \theta_0}{r(r + \delta)}(1 - e^{-rT}) - \frac{\beta e^{-rT}}{2r(r + \delta)}(\underline{E}\theta + \bar{E}\bar{\theta}) - \frac{1}{2}K(\underline{E})e^{-rT} - \frac{1}{2}K(\bar{E})e^{-rT}. \quad (24)$$

⁸ Note 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 $M_\infty = \beta E_1 / \delta$.

The values of \underline{E} and \bar{E} must be chosen optimally to maximize $W_T(\underline{E}, \bar{E})$. Setting the derivatives of $W_T(\underline{E}, \bar{E})$ with respect to \underline{E} and \bar{E} equal to zero, the optimal emission levels are:

$$\underline{E}^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\theta}{2k_2r(r+\delta)}, \quad (25)$$

$$\bar{E}^* = E_0 + \frac{k_1}{2k_2} - \frac{\beta\bar{\theta}}{2k_2r(r+\delta)}. \quad (26)$$

Should we reduce emissions now or wait until time T so that we can observe θ_T ? As before, we can compare W_0 to W_T , but now we must account for the fact that the amount of emission reduction is determined optimally at the time of adoption, i.e., at $t = 0$ or 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 Eq. (24) and E_1^* into Eq. (20) gives

$$\Delta W_T = \frac{k_1}{2k_2} \left[\frac{\beta\theta_0}{r(r+\delta)} - \frac{k_1}{2} \right] (1 - e^{-rT}) - \frac{\beta^2\theta_0^2}{4k_2r^2(r+\delta)^2} + \frac{\beta^2(\theta^2 + \bar{\theta}^2)}{8k_2r^2(r+\delta)^2} e^{-rT}. \quad (27)$$

Using Eqs. (22), (25), and (26), we can calculate that for our numerical example, $E_1^* = 191,667$ tons/yr, $\underline{E}^* = 295,833$, 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 opportunity 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 $\theta_T = \underline{\theta}$.

This numerical outcome is, of course, 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 greater reductions in E are optimal for all possible values of θ (now $E_1^* = 116,667$, $\underline{E}^* = 220,833$, and $\bar{E}^* = 12,500$), so that the sunk benefit of reducing E immediately is larger, and the sunk cost is 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 Eqs. (22), (25), and (26) again, we see that $E_1^* = 191,667$ tons/yr as before, but $\underline{E}^* = 400,000$ tons/yr, $\bar{E}^* = 0$, and $\Delta W_T = \$0.503$ billion.⁹ 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.

2.4. Summary

In this section we examined a highly simplified problem in which there are only two possible times at which a policy can be adopted—now, or a fixed time T in the future. Nonetheless, the examples illustrate how the optimal timing of policy adoption can be

⁹ Using Eq. (26), $\bar{E}^* = -16,667$. But we assume that negative values of E are not possible, so that E will be reduced to 0 if $\theta_T = \bar{\theta}$.

affected in opposing ways by the interaction of uncertainty with each of two kinds of irreversibilities. For example, by reducing the pollutant decay rate (i.e., by making environmental damage more irreversible), we increased the opportunity benefit of early policy adoption to the point where it outweighed the opportunity 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. We turn to that next.

3. A continuous-time model

When the time of adoption is a free choice variable, the problem of maximizing the present value function given by Eq. (3) becomes a classic optimal stopping problem: We must find the threshold curve, $\theta^*(M)$, that triggers policy adoption.

I generalize the model in Pindyck (2000) by allowing both θ_t and M_t to evolve stochastically. Specifically, I will assume that θ_t follows a geometric Brownian motion:

$$d\theta = \alpha\theta dt + \sigma_1\theta dz_1 \tag{28}$$

with $\alpha < r$, and that M follows a controlled arithmetic Brownian motion:

$$dM = (\beta E - \delta M) dt + \sigma_2 dz_2. \tag{29}$$

There is no reason to expect stochastic fluctuations in θ and M to be correlated, so I will assume that $\mathcal{E}_t(dz_1 dz_2) = 0$ for all t . Finally, we will work with a social benefit function that is quadratic in M , i.e.,

$$B(\theta, M) = -\theta M^2. \tag{30}$$

For simplicity, I will assume that policy adoption implies reducing emissions from E_0 to zero, at a sunk cost of $K = kE_0$. The problem is to find a rule for policy adoption that maximizes the net present value function of Eq. (3) subject to Eq. (28) for the evolution of θ , and Eq. (29) 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 $E_t = E_0$). Likewise, let $W^A(\theta, M)$ denote the value function for the “adopt” region (in which $E_t = 0$). Since $B(M_t, \theta_t) = -\theta_t M_t^2$, we know that $W^N(\theta, M)$ must satisfy the following Bellman equation:

$$rW^N = -\theta M^2 + (\beta E_0 - \delta M)W^N_M + \alpha\theta W^N_\theta + \frac{1}{2}\sigma_1^2\theta^2 W^N_{\theta\theta} + \frac{1}{2}\sigma_2^2 W^N_{MM}. \tag{31}$$

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

$$rW^A = -\theta M^2 - \delta M W^A_M + \alpha\theta W^A_\theta + \frac{1}{2}\sigma_1^2\theta^2 W^A_{\theta\theta} + \frac{1}{2}\sigma_2^2 W^A_{MM}. \tag{32}$$

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:

$$W^A(0, M) = 0, \tag{33}$$

$$W^N(0, M) = 0, \tag{34}$$

$$W^N(\theta^*(M), M) = W^A(\theta^*(M), M) - K, \quad (35)$$

$$W_\theta^N(\theta^*(M), M) = W_\theta^A(\theta^*(M), M) \quad (36)$$

and

$$W_M^N(\theta^*(M), M) = W_M^A(\theta^*(M), M). \quad (37)$$

Here, $\theta^*(M)$ is a free boundary, which must be found as part of the solution, and which separates the adopt from the no-adopt regions. It is also the solution to the stopping problem: Given M , the policy should be adopted if $\theta \geq \theta^*(M)$. Boundary conditions (33) and (34) reflect the fact that if θ is ever zero, it will remain at zero thereafter. Condition (35) is the value matching condition; it simply says that when $\theta(M) = \theta^*(M)$ and the option to adopt the policy is exercised, the payoff net of the sunk cost $K = kE_0$ is $W^A(\theta^*(M), M) - K$. Finally, conditions (36) and (37) are the “smooth pasting conditions”; if adoption at $\theta^*(M)$ is indeed optimal, the derivatives of the value function must be continuous at $\theta^*(M)$.

3.1. Obtaining a solution

Although Eq. (32) can be solved analytically, it is not possible to obtain an analytical solution for Eq. (31) and the free boundary $\theta^*(M)$. These equations can be solved numerically, although doing so is nontrivial because (31) is an elliptic partial differential equation. However, a complete analytical solution is possible if we set the decay rate, δ , to zero. Little is lost by doing so, and that is the approach I take here.

With $\delta = 0$, the analytical solution for $W^A(\theta, M)$ is

$$W^A(\theta, M) = -\frac{\theta M^2}{r - \alpha} - \frac{\sigma_2^2 \theta}{(r - \alpha)^2}. \quad (38)$$

To find a solution for $W^N(\theta, M)$, surmise that it has the form

$$W^N(\theta, M) = \theta^\gamma G(M) - \frac{\theta M^2}{r - \alpha} - \frac{2\beta^2 E_0^2 \theta}{(r - \alpha)^2} - \frac{2\beta E_0 \theta M}{(r - \alpha)^2} - \frac{\sigma_2^2 \theta}{(r - \alpha)^2}, \quad (39)$$

where $G(M)$ is an unknown function, with $G'(M) > 0$ and $G(0) > 0$. We will verify that the solution is indeed of this form. In particular, we will try solutions for which $G(M) = ae^{\eta M}$, so that the homogeneous solution to the differential equation would be of the form $W_h^N = a\theta^\gamma e^{\eta M}$. Substituting this into the differential equation, rearranging and canceling terms, gives the following equation:

$$r = \alpha\gamma + \frac{1}{2}\sigma_1^2\gamma(\gamma - 1) + \frac{1}{2}\sigma_2^2\eta^2 + \beta E_0\eta. \quad (40)$$

Note that this is an equation in both γ and η . Hence, we cannot solve this without making use of other boundary conditions. In addition, we need to find the value of a .

In total, there are four unknowns for which solutions must be found: γ , η , a , and $\theta^*(M)$. To solve for these four unknowns, we make use of Eq. (40), along with

boundary conditions (35)–(37).¹⁰ Boundary condition (35) implies

$$-\frac{2\beta^2 E_0^2 \theta^*}{(r-\alpha)^3} - \frac{2\beta E_0 \theta^* M}{(r-\alpha)^2} + a(\theta^*)^\gamma e^{\eta M} = -K, \tag{41}$$

boundary condition (36) yields

$$-\frac{2\beta^2 E_0^2}{(r-\alpha)^3} - \frac{2\beta E_0 M}{(r-\alpha)^2} + a\gamma(\theta^*)^{\gamma-1} e^{\eta M} = 0, \tag{42}$$

and boundary condition (37) yields

$$a\eta(\theta^*)^\gamma e^{\eta M} = -\frac{2\beta E_0 \theta^*}{(r-\alpha)^2}. \tag{43}$$

Eqs. (40)–(43) are all nonlinear, and must be solved simultaneously for γ , η , a , and $\theta^*(M)$. This is most easily done by multiplying Eq. (42) by θ^* , and then using Eq. (43) to eliminate $a(\theta^*)^\gamma e^{\eta M}$ from that equation and from Eq. (41). The remaining three equations then yield the following solution. Defining $\Omega(M) \equiv (r-\alpha)M + \beta E_0$, the exponent $\gamma(M)$ is given by

$$\gamma = \frac{\alpha\Omega^2 - \frac{1}{2}\sigma_1^2\Omega^2 + (r-\alpha)\beta E_0\Omega}{\sigma_1^2\Omega^2 + \sigma_2^2(r-\alpha)^2} \left[-1 + \sqrt{1 + \frac{2r[\sigma_1^2\Omega^2 + \sigma_2^2(r-\alpha)^2]}{[\alpha\Omega - \frac{1}{2}\sigma_1^2\Omega + (r-\alpha)\beta E_0]^2}} \right], \tag{44}$$

the exponent $\eta(M)$ is given by

$$\eta = (r-\alpha)\gamma/\Omega \tag{45}$$

and the optimal stopping boundary is given by

$$\theta^*(M) = \frac{\gamma(r-\alpha)^3 K}{2(\gamma-1)\beta E_0 \Omega(M)}. \tag{46}$$

Finally, the variable a is given by

$$a = \frac{2\beta E_0}{(r-\alpha)^2} \eta(\theta^*)^{1-\gamma} e^{-\eta M}. \tag{47}$$

Eqs. (44) and (46) completely determine the solution to the optimal timing problem: Emissions should be reduced to zero when $\theta \geq \theta^*(M)$. It can be shown that $\theta^*(M)$ is a declining function of M , as we would expect: the greater is the current stock of pollutant, the lower is the critical social cost variable, θ^* , at which the emission-reducing policy should be adopted.¹¹ These equations, together with Eqs. (45)

¹⁰ Eq. (40) is a quadratic in γ , so condition (34) is used to rule out one of the two solutions for γ .

¹¹ Note that

$$\frac{d\theta^*}{dM} = \frac{\partial\theta^*}{\partial\gamma} \frac{\partial\gamma}{\partial\Omega} \frac{\partial\Omega}{\partial M} + \frac{\partial\theta^*}{\partial\Omega} \frac{\partial\Omega}{\partial M}.$$

From the equations above, $\partial\Omega/\partial M = r-\alpha > 0$ and $\partial\theta^*/\partial\gamma < 0$. With some messy algebra, it can be shown that $\partial\gamma/\partial\Omega > 0$, and thus $d\theta^*/dM < 0$.

and (47) also determine the value of the option to adopt the emission-reducing policy, namely $a\theta^\gamma e^{\eta M}$.

3.2. Characteristics of the solution

By calculating solutions for different combinations of values for the parameters σ_1 and σ_2 , we can explore how economic and ecological uncertainties affect the optimal timing of policy adoption. To do this, we must choose a range of values for these parameters, as well as values for the other parameters in the model, that are consistent with pollution and cost levels that could arise in practice. Although these calculations are meant to be largely illustrative, they will be done in the context of GHG emissions and global warming.

For the real interest rate, absorption parameter, and initial level of emissions we will use the same values as in the two-period model: $r=0.04$, $\beta=1$, and $E_0=300,000$ tons/yr. With the pollutant decay rate, δ , equal to zero, this rate of emissions would add 30 million tons to the pollutant stock after 100 yr.¹² We will consider current pollutant stocks (of human origin) in the range of 10–150 million tons. We will set the present value of the cost of policy adoption, K , at \$4 billion; although the actual cost is likely to be much larger, over a long period of time, much of it should be reversible. We will initially set α , the expected percentage rate of growth of θ , to zero, although we will also calculate solutions for $\alpha=0.01$.

Finally, as initial values for the volatility parameters, we use $\sigma_1=0.2$ and $\sigma_2=1,000,000$, although we will also vary these numbers. This value for σ_1 implies an annual standard deviation of 20 percent for the social cost generated by the pollutant stock, and a standard deviation of 200 percent for a 100-yr time horizon, a number that is consistent with current uncertainties over this cost. The value for σ_2 implies a standard deviation of 10 million tons for the stock level after 100 yr, which is one-third of the expected increase in the stock from unabated emissions. (In the case of climate change, M is usually viewed as largely nonstochastic—see, e.g., Nordhaus (1994)—and the real uncertainty is over *temperature*, which is the cause of economic damage. However, changes in temperature depend, with a lag, on changes in M , so M is a good proxy for the stochastic state variable.)

Fig. 1 shows the critical threshold $\theta^*(M)$ for values of M ranging from 0 to 16 million tons. The middle curve is $\theta^*(M)$ for the base values of $\sigma_1=0.2$ and $\sigma_2=1,000,000$, and $\theta^*(M)$ is also shown for $\sigma_1=0$, $\sigma_2=1,000,000$ and $\sigma_1=0.4$, $\sigma_2=2,000,000$. Note that these curves are downward sloping, as we would expect—a larger M implies a larger social cost, and thus a lower value of θ at which it is optimal to adopt the policy.

For these parameters, the value of waiting is large. To see this, we can calculate a traditional net present value for the adoption decision at the critical threshold $\theta^*(M)$. Fig. 2 shows (for each of the three cases in Fig. 1) the present value of the gains

¹² Setting $\delta=0$ is a reasonable approximation for GHGs—the actual decay rate has been estimated to be 0.5 percent or less.

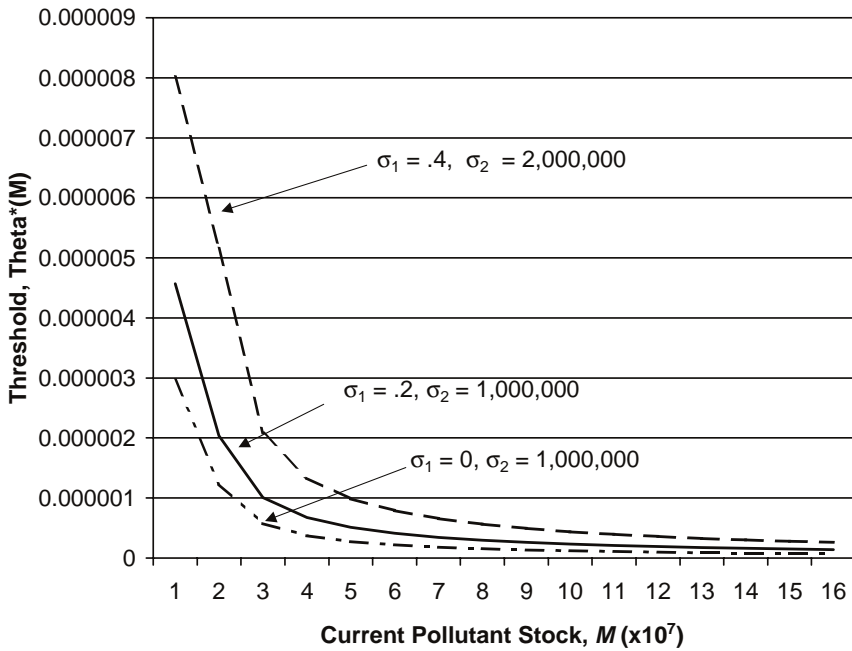


Fig. 1. Critical threshold, $\theta^*(M)$.

from policy adoption relative to the cost of adoption, K . Note that from Eqs. (38) and (39), this ratio is given by

$$PV/K = [2\beta^2 E_0^2 \theta / (r - \alpha)^3 + 2\beta E_0 \theta M / (r - \alpha)^2] / K. \tag{48}$$

Under a traditional NPV rule, adoption would occur when this ratio exceeds one. Observe from Fig. 2, however, that for small values of M policy adoption is optimal only when this ratio is considerably greater than one, and for our base case values of σ_1 and σ_2 this ratio exceeds two for all values of M in the range considered.

Observe from Figs. 1 and 2 that $\theta^*(M)$ and the ratio PV/K flatten out once K exceeds 4 or 5. The reason is that when M is large, continued emissions makes little difference for uncertainty over future values of M , because they contribute little in percentage terms to the expectations of those future values. (Recall from Eq. (29) that M follows a controlled arithmetic Brownian motion). Thus for large M , the volatility of M , i.e., σ_2 , makes a negligible contribution to the value of waiting. This can be seen from the bottom curve in Fig. 2, for which $\sigma_1 = 0$. For large M the ratio PV/K is only slightly > 1 . When $\sigma_1 > 0$, the ratio exceeds one, but only because of uncertainty over the future value of θ and hence the future social cost of added emissions.

This illustrates an important difference between the effects of economic versus ecological uncertainty. If stochastic fluctuations in the pollutant stock are arithmetic in nature, they would create uncertainty over the future social cost of continued emissions only because the social benefit function $B(\theta, M)$ is quadratic in M . Stochastic

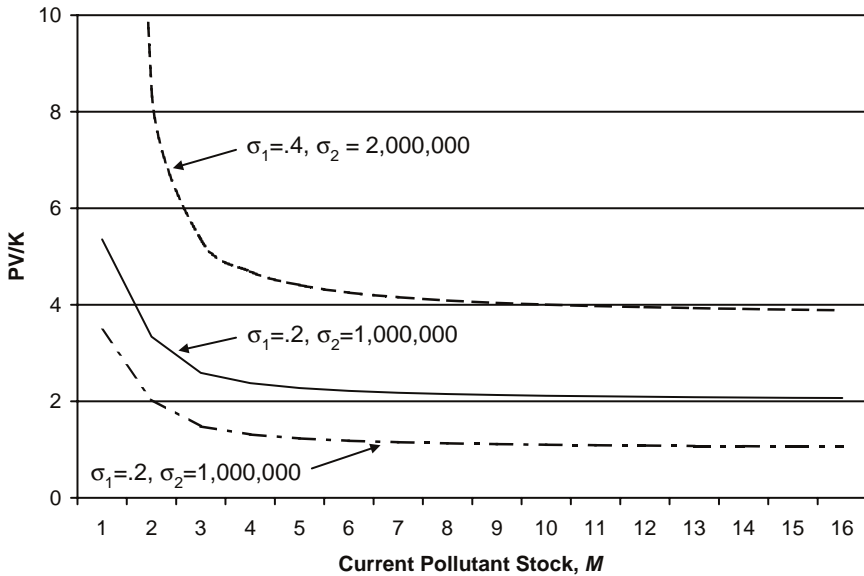


Fig. 2. Traditional present value comparison. (Shows present value of benefits from immediate adoption relative to cost, K , at critical threshold $\theta^*(M)$.)

fluctuations in the economic cost variable θ , however, shift the entire social benefit function for every level of M . Of course, one might argue that the process for M should be modelled as a controlled geometric Brownian motion, so that the last term in Eq. (29) is $\sigma_2 M dz_2$. I have seen little empirical support for this, however, and one would expect that unpredictable increases or decreases in M are due largely to under- or overpredictions of emissions levels from various sources, and thus should not depend on the overall level of the pollutant stock.¹³

Fig. 3 shows the critical threshold $\theta^*(M)$ as a function of σ_1 for a value of M equal to 50 million tons, and for the drift parameter α set at zero and at 0.01. As with models of irreversible investment, increases in uncertainty over the future “payoffs” from reduced emissions increase the value of waiting, and raise the critical threshold $\theta^*(M)$. Increasing the drift parameter, α , from 0 to 0.01 reduces the threshold at each value of M ; a higher value of α implies higher expected future payoffs from reducing emissions now.

Fig. 4 shows $\theta^*(M)$ as a function of σ_2 , the volatility of M , again for a value of M equal to 50 million tons, and for α equal to 0 and 0.01. The threshold $\theta^*(M)$ increases with σ_2 , but only slowly. As discussed above, with $M=50$ million, continued emissions increase M by a small amount in percentage terms over a 20- or 30-yr period, so that stochastic fluctuations in M can have only a small effect on the value of waiting (and

¹³ As mentioned earlier, in the case of global warming most of the uncertainty is over the future temperature, which in turn depends on M . However, there is little reason to expect stochastic changes in temperature to depend on the current temperature level.




Fig. 3. Dependence of critical threshold, θ , on σ_1 ($\sigma_2 = 1, \alpha = 0$ and 0.01).




Fig. 4. Dependence of critical threshold, θ , on σ_2 ($\sigma_1 = 0.2, \alpha = 0$ and 0.01).

that effect is due to the convexity of $B(\theta, M)$). Thus changes in σ_2 can have only a small effect on the threshold that triggers policy adoption. (But note that changes in σ_2 will have a larger effect on the threshold if M is small.) A change in α , however, will again have a large effect on the threshold because it changes the expected future payoffs from emissions reductions.

4. Conclusions

Environmental policies, which impose sunk costs on society, are often adopted in the face of considerable uncertainties over the flow of net benefits that they will generate. On the other hand, the adoption of those policies also yields “sunk benefits” in the form of averted irreversible environmental damage. These opposing incentives for early versus late adoption were illustrated in the context of a simple two-period model in an emissions-reducing policy that could be adopted either now or at some fixed time in the future. This timing problem was explored again through the use of a continuous-time model in which adoption could occur at any time, and there is uncertainty over the future economic benefits of policy adoption, and over the future evolution of the pollutant stock.

In both cases, I focused largely on a one-time adoption of an emission-reducing policy. One might argue that policies could instead be adopted or changed on an “incremental” basis; for example, a carbon tax could be imposed and then adjusted every few years in response to the arrival of new information regarding global warming and its costs. In reality, however, policy adoption involves large sunk costs of a political nature—it is difficult to adopt a new policy in the first place, or to change one that is already in place.

In addition, I assumed that policy-induced costs were completely sunk, and that policy adoption is irreversible in that the policy could not be partially or totally reversed in the future. (In Section 2.2, however, I examined the implications of allowing for a single policy reversal.) 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 may be extreme. Richer models are needed to explore the implications of relaxing this assumption.

Finally, one could argue that my specification of the stochastic process for the stock of pollutant, M , is restrictive. This process could easily be generalized, but it would then be necessary to obtain numerical solutions of the differential equations for the value functions. That would be a logical extension of this work, because one could then also allow for a nonzero decay rate, δ .

Acknowledgements

This paper was written while the author was a Visiting Professor at the Harvard Business School, and the hospitality of that institution is gratefully acknowledged.

I am also grateful to M.I.T.'s Center for Energy and Environmental Policy Research for its financial support of the research underlying this paper.

References

- Arrow, K.J., Fisher, A.C., 1974. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics* 88, 312–319.
- Bernanke, B.S., 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* 98, 85–106.
- Cline, W.R., 1992. *The Economics of Global Warming*. Institute for International Economics, Washington, DC.
- Conrad, J.M., 1992. Stopping rules and the control of stock pollutants. *Natural Resource Modeling* 6, 315–327.
- Dixit, A., Pindyck, R.S., 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.
- Hammitt, J.K., Lempert, R.J., Schlesinger, M.E., 1992. A sequential-decision strategy for abating climate change. *Nature* 357, 315–318.
- Hendricks, D., 1992. Optimal policy response to an uncertain threat: the case of global warming. Unpublished manuscript, Kennedy School of Government, Harvard University, March 1992.
- Henry, C., 1974. Investment decisions under uncertainty: the irreversibility effect. *American Economic Review* 64, 1006–1012.
- Kelly, D.L., Kolstad, C.D., 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23, 491–518.
- Kendrick, D., 1981. *Stochastic Control for Economic Models*. McGraw-Hill, New York.
- 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 University 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 1998.
- Nordhaus, W.D., 1994. *Managing the Global Commons: The Economics of Climate Change*. MIT Press, Cambridge, MA.
- Peck, S., Teisburg, T., 1992. Global warming uncertainties and the value of information: an analysis using CETA. *Resource and Energy Economics* 15.
- Pindyck, R.S., 1996. Sunk costs and sunk benefits in environmental policy. Unpublished Working Paper, M.I.T. Center for Energy and Environmental Policy, April 1996.
- Pindyck, R.S., 2000. Irreversibilities and the timing of environmental policy. *Resource and Energy Economics* 22, 233–259.
- 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.
- 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.