Timing and Scope of Emission Reductions for Airborn Particulate Matter: A Simplified Model

Michael J. Phelan

NRCSE

Technical Report Series

NRCSE-TRS No. 029

The **NRCSE** was established in 1997 through a cooperative agreement with the United States Environmental Protection Agency which provides the Center's primary funding.

TIMING AND SCOPE OF EMISSION REDUCTIONS FOR AIRBORNE PARTICULATE MATTER: A SIMPLIFIED MODEL

Michael J Phelan[1]

Chapman University

Abstract

Environmental health policy decisions are characterized by irreversibility and uncertainty of an economic, ecological and biomedical nature. For this reason decision-makers may choose to exercise when possible some discretion over the timing and scope of policy. Problems of this kind fall within the framework of the theory of irreversible investments as applied to the sunk costs and sunk benefits of environmental regulation. This work is an application of the basic theory to the problem of timing and scope of emissions reductions for airborne particulate matter. Particular attention is given to the representation of health effects in a simplified model of benefit-cost uncertainty. The results describe analytically and illustrate numerically a decision rule for optimal policy. The exposition is designed primarily for the benefit of quantitative policy analysts in the areas of environmental and health regulation.

Keywords: Environmental health regulation, airborne particulate matter, benefit-cost analysis, irreversible investments, stochastic control

Introduction

There are growing national concerns about the costs, reach, and effectiveness of environmental, health and safety regulation. The push is toward legislative reforms favoring greater reliance on economic analysis in policy decisions. To discuss this issue a group of prominent economists met recently under the auspices of the American Enterprise Institute, the Annapolis Center, and Resources for the Future. As they emerged from these discussions, Arrow and co-authors (1996) articulated a set of principles for guiding and improving quality in the use of benefit-cost analysis in environmental, health and safety regulation. Of interest

[1]Although the research described in this article has been funded in part by the United States Environmental Protection Agency through agreement CR25173-01-0 to the University of Washington, it has not been subjected to the Agency's required peer and policy review and therefore does not necessarily reflect the views of the Agency and no official endorsement should be inferred.

here was mention of the need to account for irreversibility of decisions and, repeatedly, the ubiquity of uncertainty.

The present work describes an application of the basic theory of irreversible investments under uncertainty to the problem of regulating airborne particulate matter. Policy is defined as a solution to a stochastic control problem for the optimal timing and scope of emissions reduction. Familiar readers will recognize that such problems can easily become intractable, except to deeper mathematics or intensive numerical methods. In the interest of keeping both the theory and the practice accessible, however, the problem is considered here only in the context of a simplified model of benefit-cost uncertainty. Optimal policy is then derived explicitly and studied numerically in a sensitivity analysis across selected decisional elements. This project is designed for the benefit of quantitative policy analysts to serve as an introduction to the framework's potential to improve decision-making when setting air quality standards.[2]

The U.S. Environmental Protection Agency's (EPA) new National Ambient Air Quality Standards for airborne particulate matter of July 1997 underscores the growing concern about the effects of particulate matter on environment and human health. The present analysis is an opportunity to evaluate the practical importance of this theory and to determine whether the sunk costs of policy adoption outweigh the sunk benefits of environmental preservation. In the case of the EPA's new standards for particulate matter, a substantial flow of benefits are anticipated to be derived directly from the positive health effects of improved air quality. Health effects are therefore represented explicitly in the present model, rather than implicitly as is often done in applications of the theory to the economic analysis of global warming.

The proposed theory can be put in several ways, but perhaps the most accessible for our purposes is its expression by Pindyck (1996) and the formal treatment here owes its essential features to his basic theory. Starting from an analogy with irreversible investments, environmental policy decisions are framed in terms of problems in stochastic optimization and control. The general framework includes problems of optimal timing of policy and se-

[2]An exception to this expository standard occurs in the technical demonstration of the main result, where reference is made to the analytical framework for economic decision-making of Brekke and Øksendal (1991, 1994). Their theory is broadly applicable to the present problem and to more advanced applications.

quential or incremental policy response. It also provides for a natural analysis of the role of uncertainty and learning in the design and implementation of policy. Though illustrated particularly in terms of economic and ecological factors in the context of global CO_2 emissions policy, the theory is posed generally enough to cover a broad range of environmental problems.

A brief history of this approach might begin with a discussion in Weisbrod (1964) about the 'option value' of preserving national parks for later generations who may want to visit pristine lands. Further elaborations emerged with more than one interpretation of Weisbrod's notion, but the branch followed here is perhaps best represented analytically by the contributions of Fisher and co-authors (1972), Arrow and Fisher (1974), Henry (1974), Fisher and Krutilla (1985), Hanemann (1989) and finally Pindyck (1996). Their research issued the development of a statistical theory of optimal decision making in environmental economics and management. The driving force in this theory is irreversibility. Uncertainty plays a qualifying but inescapable role. The result today is that, at least in theory, environmental policy decisions must reconcile certain opportunity costs and benefits, the value of flexible timing or scope of regulation, and the prospect that the future will in some way bring missing or incomplete information to light.

The full weight of the theory is beyond the scope and purpose of the present work. A more sophisticated model would substantially change the results, so only qualitative implications should be drawn from the present model and the numerical illustrations. At the same time, though the model highly stylized, the results from the present model may nevertheless serve as a point of reference, in the manner suggested by Cox (1997), for comparison with the results of numerical simulations of more detailed models.

Benefit-cost uncertainty: A simplified model

This section describes a simplified model for economic analysis of timing and scope of regulation for airborne particulate matter. The model deals only with economic or benefit-cost uncertainty; otherwise, the stock of pollutants and the stock of loss-of-health are known with certainty. That the latter kind of uncertainty plays no role here is a consequence of the linear structure of the social costs and the nature of the payoff from regulation. Alternative models are discussed later as the focus of continuing research.

Readers familiar with Pindyck's work on global warming will recognize the model. It is one among several models considered in his work for which there exists an explicit analytical solution, a solution which may be readily examined for its sensitivity to the key elements of decision. The purpose of this discussion is to introduce the decision-making framework for environmental regulation in the context of an environmental *health* regulation where the regulatory payoff may also be derived from the anticipated benefits to the public health. In the case of regulating airborne particulate matter, this means reduced rates of upper respiratory diseases like asthma. A new feature in the present model, therefore, is the representation of health factors in addition to ecological factors among the decisional elements defining policy. The results will also examine the effects of these factors on the distribution of the optimal timing of policy.

First, suppose that pollution enters the environment at the equilibrium rate ϵ and decays at the rate δ. In a world of certainty the stock of pollutant converges in the long run to the ratio ϵ/δ. Second, suppose conceptually that there is an underlying stock of loss of human health and that $\tilde{\eta}$ is the rate of incidence or production of loss-of-health while $\tilde{\gamma}$ is the rate of loss-of-health decay. The stock of loss-of-health converges in the long run to the ratio $\tilde{\eta}/\tilde{\gamma}$. Next, suppose that society chooses to adopt a policy at calendar time T with the effect of reducing the emissions rate by $100u\%$, for $0 < u \leq 1$. In this case, suppose that the stock of pollution, M, follows the trajectory given by:

$$M_t = \begin{cases} \frac{\epsilon}{\delta}(1 - \mathrm{e}^{-\delta t}) + M_0 \mathrm{e}^{-\delta t} & \text{for } t < T \\ \frac{(1-u)\epsilon}{\delta}(1 - \mathrm{e}^{-\delta t}) + M_T \mathrm{e}^{-\delta t} & \text{for } t \geq T. \end{cases} \tag{1}$$

Suppose further that as policy takes effect a marginal increase in the optimal health trajectory is anticipated in response to the positive health effects associated with a drop in pollution. In particular, suppose that the production rate decreases to $\tilde{\eta}\mathrm{e}^{-du}$, for some positive elasticity d. In this case, suppose that the stock of loss-of-health, $\tilde{H}$, follows the trajectory given by:

$$\tilde{H}_t = \begin{cases} \frac{\tilde{\eta}}{\tilde{\gamma}}(1 - \mathrm{e}^{-\tilde{\gamma} t}) + \tilde{H}_0 \mathrm{e}^{-\tilde{\gamma} t} & \text{for } t < T \\ \frac{\tilde{\eta}\mathrm{e}^{-du}}{\tilde{\gamma}}(1 - \mathrm{e}^{-\tilde{\gamma} t}) + \tilde{H}_T \mathrm{e}^{-\tilde{\gamma} t} & \text{for } t \geq T. \end{cases} \tag{2}$$

Policy then effects a downward shift in both the equilibrium level of the stock of airborne particulate matter and the stock of human loss-of-health. The presence of knock-on health effects is particularly important to benefit-cost analysis of policy for airborne particulate matter.[3]

For the purpose of economic analysis, let A denote the unit cost of emissions in monetary terms as given by the Itô equation:

$$(3) \qquad dA_t = \alpha A_t dt + \sigma^a A_t W(dt), \qquad \text{or} \qquad A_t = A_0 \mathrm{e}^{\alpha t - \frac{1}{2}(\sigma^a)^2 t + \sigma^a W_t}$$

where W is a Wiener process and α and σ^a are fixed parameters. As a geometric Brownian motion with appreciate rate α and volatility σ^a, unit cost of emissions is the only source of benefit-cost uncertainty in the problem. Suppose that the flow of social costs of pollution, $X^{u,T}$, is given by the linear expression:

$$(4) \qquad X_t^{u,T} = -A_t(M_t + \lambda \tilde{H}_t),$$

where λ denotes the cost basis for a unit of loss-of-health per unit cost emissions. This is consistent with one of the models for benefit-cost analysis introduced and developed by Pindyck (1996). The new feature in the present model is the explicit abstraction of health effects. The idea is that A responds to tastes and preferences, while $M + \lambda \tilde{H}$ expresses damages due to pollution in the environment and human loss-of-health, the chief externalities of economic activity in the model. The basic idea for this expression is derived from the conceptual framework for economic analysis in Nordhaus (1991), where consumptive preferences arise from a utility function and the damages from the disutility of pollution.[4]

The model is interpreted to say that society derives the fruits of economic activity in part by burning fossil fuels and so on. The economic process, A, responds to shifts in tastes,

[3] Grossman (1972) first developed the notion of health capital, H say, as associated with its productive element 'healthy time.' There is then a parameter, Ω_0 say, representing the maximal available capital, so that $\Omega_0 - H$ expresses the present notion of loss-of-health stock. In wage studies in econometrics, for example the work of Cropper (1977, 1981), there are 365 healthy days available so that loss-of-health amounts to the number of unhealthy or sick days without work.

[4] This conception, as pointed out in Pindyck (1996), has enormous potential for the purpose of generalizing the approach. Directions for this are discussed later in this work.

technologies, and patterns of resource and health-care utilization. Its appreciation rate, α, is that of the underlying rate of growth of consumption in the economy. Chief among the externalities in the model is the stock of pollutants, M, which accumulates with economic activity at rate ϵ and decays at rate δ. The parameter δ specifies the ambient rate of removal by the environment, so a strictly positive removal rate indicates that the effect of pollution on the ecology is at least partially reversible. Next, the knock-on health effects of pollution are at least partially responsible for the production of a loss stock of human health, $\tilde{H}$, which accumulates at rate $\tilde{\eta}$ with the incidence of mortality, morbidity, and respiratory disease like asthma. At the same time, with gains in health capital, the stock of loss-of-health decays at the underlying rate $\tilde{\gamma}$. The model says that only the production of lost health, and not its decay, depends on the emissions rate of airborne particulate matter. In this case, the control of emissions has the potential to modify the trajectories of the loss-of-health stock by permanently lowering the rate of production of emissions-related illness and disease.[5]

A policy that reduces stock levels of pollutant by δ_M say incurs a benefit given by $A\delta_M$ in ecological terms and, assuming a knock-on reduction $\delta_{\tilde{H}}(\delta_M)$ in lost health stock, a health benefit given by $\lambda A\delta_{\tilde{H}}(\delta_M)$, where λ is the cost basis for a unit of loss-of-health per unit cost emissions. The net benefit of such a policy is then given by the present value of the flow of social benefit, $A(\delta_{\tilde{H}}(\delta_M) + \delta_M)$, minus the adoption cost of the policy. The tradeoff here is then between the costs of the reduction and society's willingness to pay for the benefits, reflecting both the value placed on income or consumption and the anticipated improvements in public health.

The policy problem is to devise a strategy or control for regulating the rate of emissions of airborne particulate matter. According to the model, each control may be represented by a pair (u, T), say, where T is a stopping time with respect to the filtration or history generated by the economic process A and u is a level of reduction in $[0, 1]$. The implied regulation typically entails a commitment to an adoption cost amounting to the (random) variable $K^{u,T}$ or simply K^u when such costs depend only on the level or scope of the reduction. Unless policy is reversible, the adoption cost is completely sunk as for example would be the

[5]While it may also be natural to expect a pollution effect on the rate of loss-of-health decay, such an effect significantly complicates the analysis of the model beyond the scope and intention of the present work.

cost of retrofitting a production facility with an emissions control device. Suppose finally that the social value of time is given by the discount rate r, which strictly exceeds the appreciation rate, α, of unit cost of emissions. In this case, the net present value of policy is given by

$$V^{u,T}(A_0, \tilde{H}_0, M_0) = I\!E_0 \int_0^\infty X_t^{u,T} \mathrm{e}^{-rt} dt - I\!E_0 K^u \mathrm{e}^{-rT}, \tag{5}$$

expressing the monetary value associated with a particular emissions policy, that is (u, T) from the class of admissible controls. Notice that $I\!E_0$ is the conditional expectation given the initial data $(A_0, \tilde{H}_0, M_0)$. The value, V^*, of an optimal policy is then given by

$$V^*(A_0, \tilde{H}_0, M_0) = \sup_{u,T} [I\!E_0 \int_0^\infty X_t^{u,T} \mathrm{e}^{-rt} dt - I\!E_0 K^u \mathrm{e}^{-rT}], \tag{6}$$

as an expression of the monetary value of an optimal policy for the control of emissions of particulate matter.

Ordinarily, the optimal control, (u^*, T^o) say, at which maximal value is attained, is difficult to characterize explicitly. Not so for the simplified model of benefit-cost uncertainty being considered here. The discussion therefore will focus on characterizing the solution in terms of optimal timing of optimal levels of reduction and on providing numerical illustrations of the main result.

Elements of decision and optimal policy in a world of certainty

This section contrasts the present model with a traditional benefit-cost analysis. The contrast exposes the key elements of decision-making in the formulation of optimal policy. The phrases "sunk costs" and "sunk benefits" express a central concept in the analysis of irreversible decisions. A sunk cost is the irrecoverable cost of adopting an environmental policy, such as the cost of installing new scrubbers on factories, scrapping old machinery for new fuel-efficient models, or paying higher prices for better-grade fuels. A sunk benefit, in contrast, is a negative opportunity cost or preventive benefit of adopting an environmental policy, such as the benefit of avoiding irreversible environmental damage, of preserving fragile ecosystems, of saving human lives, or of reducing morbidity. Understanding the trade-offs imposed by such costs and benefits is key to understanding a modern benefit-cost analysis.[6]

[6]This section is again designed for the benefit of the quantitative policy analyst who may not be entirely

The basic themes are demonstrated here in a world of certainty. This is achieved in the present model by removing the economic uncertainty from the picture, namely taking σ^a equal 0 so that the economic process, A, is given by $t \to A_0 e^{\alpha t}$. A traditional benefit-cost analysis compares the value of adopting policy now ($T = 0$) with the value of never adopting policy ($T = \infty$). For this purpose, let r denote the social discount rate, where r exceeds α. The net present value, NPV_0, of the proposed regulation is then given by:

$$\text{(7)} \qquad \mathrm{NPV}_0 = \frac{\lambda A_0(\tilde{\eta} - \tilde{\eta}e^{-du})}{(r-\alpha)(r+\tilde{\gamma}-\alpha)} + \frac{A_0 u \epsilon}{(r-\alpha)(r+\delta-\alpha)} - K^u$$

where K^u denotes the sunk cost of adopting the policy now. The leading term on the right-hand-side of Equation 7 represents the present value of the anticipated benefits of mitigating adverse effects of pollution on the public health. The second term represents the present value of the ecological benefit of reduced emissions to the environment. Society adopts the policy whenever NPV_0 is positive, meaning the benefits of pollution abatement now exceed the sunk cost of adoption. In this case, the scope of policy or level of response, u, is derived by equating the marginal cost of adoption to the marginal benefit. If $u \to K^u$ where linear, for example, the optimal scope would be 100%.

In addition to the adoption cost, K^u, Equation 7 introduces two other key elements of decision-making for environmental health policy: the terms $\mathcal{E}$ and $\tilde{\mathcal{H}}$ as given by the expressions

$$\text{(8)} \qquad \mathcal{E} = \frac{u\epsilon}{(r-\alpha)(r+\delta-\alpha)}, \qquad \text{and} \qquad \tilde{\mathcal{H}} = \frac{(\tilde{\eta} - \tilde{\eta}e^{-du})}{(r-\alpha)(r+\tilde{\gamma}-\alpha)}.$$

The ecological factor, $\mathcal{E}$, figures prominently in the work of Pindyck (1996). The health factor, $\tilde{\mathcal{H}}$, appears here because of the explicit abstraction of health effects of pollution on the flow of social costs. Expressed in monetary terms, these elements of decision give rise to the value of the ecological benefit of policy, namely $A_0\mathcal{E}$, and the value of the health benefits of policy, namely $A_0\lambda\tilde{\mathcal{H}}$.

---

familiar with the notion of irreversibility and its consequences for environmental policy. Readers familiar with these ideas may choose to skip ahead. The analogy between irreversible investments and environmental policy is already familiar to readers of the literature on global warming; see for example Birge and Rosa (1995), Conrad (1997), Nordhaus (1991), and especially Pindyck (1996) from this large and fast moving literature. Readers interested in an entirely nontechnical discussion of these ideas for the benefit of environmental health professionals may wish to consult Phelan (1998).

Evidently the value of ecological plus health benefits provides an incentive for society to act now in the face of the sunk costs of adoption. The value of the incentive decreases in ecological terms with a rising removal rate, δ, and in health terms with a rising decay rate, $\tilde{\gamma}$, indicating in both cases the partial reversibility of damages. On the other hand recall that the social costs of pollution are appreciating at rate α. The effects of irreversibility then imply, even in a world of certainty, that it may be optimal for society to delay policy until such time as discounting lowers sufficiently the present value of the sunk costs of adoption.

In particular, suppose society contemplates delaying policy until time τ, representing some fixed calendar date in future. The net present value of policy that waits until time τ rather than adopt policy now is given by:

$$\text{NPV}_\tau = K^u(1 - \mathrm{e}^{-r\tau}) - A_0(\mathcal{E} + \lambda\tilde{\mathcal{H}})(1 - \mathrm{e}^{-(r-\alpha)\tau}) \tag{9}$$

where the factors $\mathcal{E}$ and $\tilde{\mathcal{H}}$ are those of Equation 8. Society will delay policy whenever NPV_τ is positive, which depends on the interplay between the opportunity costs and benefits of doing so. Notice that the sunk cost of adoption, K^u, will always provide an incentive to wait. This is measured by the savings brought about by the difference between adoption costs today and the present value of adoption costs tomorrow, namely the opportunity cost of early exercise of society's option to adopt policy. In contrast, ecological benefits and health benefits present a general disincentive to wait in terms of the negative opportunity costs of delaying sunk benefits.

The exhibit entitled "A numerical illustration" treats these decisional elements in terms of a hypothetical scenario. The discussion illustrates tradeoffs between sunk costs and sunk benefits that are typical of environmental problems, reflecting the role particularly of irreversibility—a role further qualified by uncertainty—which significantly shifts the grounds for decision making and the ultimate design of policy.[7] Responding to problems today or

[7] A basic insight is that the omission of sunk benefits from traditional benefit-cost analysis biases now-or-never decisions against adoption of policy, while the omission of sunk costs biases now-or-never decisions toward adoption of the policy. The timing of policy on the other hand involves an interplay between the negative opportunity costs of delaying sunk benefits, since waiting entails the realization of ecological damage to the environment and adverse effects to the public health, and the depreciating present-value of adoption costs.

A numerical illustration

This numerical illustration of the elements of decision is not an attempt at an actual analysis of policy, nor is there pretense to justify the particulars on objective grounds.The values chosen for illustration may be more-or-less consistent with figures from various sources.

First, the rate of emissions, ϵ, of particulate matter (PM) is given by 10 megatons per annum. The initial social cost, A_0, is \$25 per ton per year, which cost appreciates at α equal 2.5% per year. The emission figure includes about 6.9 megatons of direct PM plus a very modest indirect contribution from the annual emissions of selected precursors among SO_2, NO_X and various volatile organic compounds. The removal rate, δ, is taken to be a fairly efficient 0.045. The equilibrium level of stock is then about 222 megatons.

Next, the initial social cost, λA_0, per 'unit loss-of-health capital' is \$5 million per unit loss-of-health per year, which cost then appreciates at 2.5% per year. The parameter d for health effects is taken to be 0.7 per unit of emissions reduction. The rate of loss-of-health decay, $\tilde{\gamma}$, is taken to be 0.025, and the rate of loss-of-health production, $\tilde{\eta}$, 6.5 units per annum. This means, for example, that the equilibrium is initially 260 units of loss-of-heath capital.

There is an open problem here of how best to calibrate the unit cost of emissions for the human loss-of-health capital. For the present purposes, the social cost of a unit of loss-of-health scales essentially, though not literally, to about 22 statistical years of life saved at \$229,000 per year. This conception of loss-of-health capital is an imperfect composite folding together various categories of human suffering in a diverse mixture of people and costs. The 50% reduction contemplated below yields about 1,700 of these 'year-units' per annum with a present-value benefit of \$16 billion.

A second, no less sticky problem lies with the notoriously elusive health effects of PM, which are difficult to establish, tend to be relatively small and appertain to large numbers of individuals. The choices made here suggest a moderate effect on loss-of-health production. Some of this moderate effect might be obtained by virtue of a cleaner environment, even in the absence of demonstrable effects of particulate matter on either the incidence or severity of respiratory disease. In response to the cleaner environment, some argue, greater public participation in outdoor recreation might incur the health benefits of a more active population.

Suppose that society contemplates halving the rate of emissions from its current rate of 10 megatons to 5 megatons per annum. The cost of adoption, $K^{0.5}$, is given by \$70.5 billion. For a social discount rate, r, of 4% per year, the decisional elements $A_0\mathcal{E}$ and $A_0\lambda\tilde{\mathcal{H}}$ are given by \$139 billion and \$16 billion, respectively, as derived from Equations 8. According to Equation 7, the net present value, NPV_0, is given by:

$$\mathrm{NPV}_0 = \$141\mathrm{b},$$

so conventional benefit-cost analysis would suggest that immediate control at 50% reduction is better than no regulation at all. On the other hand, according to Equation 9, the net present value of waiting 6 years, NPV_6, is given by:

$$\mathrm{NPV}_5 = \$1.7\mathrm{b},$$

so there is greater value in waiting. In fact, according to these data, the optimal timing for a 50% reduction of emissions is about 7.7 years. By contrast, the optimal timing for a 40% reduction turns out to be 2.5 years, so the value of waiting 6 years, or 3.5 years too late, turns out to be negative.

at full levels incurs certain opportunity costs that put value on waiting for new information or employing flexible approaches. Recognizing these effects leads the analysis toward the optimal policy in terms of timing and scope of regulation.[8]

Optimal timing in a world of certainty involves maximizing the value function of Equation 5, specifically maximizing over T for fixed u. The result here is that the optimal time, T^{O}, for the adoption of policy in a world of certainty is given by:

$$T^{\mathrm{O}} = \frac{1}{\alpha} \ln\left[\frac{rK^u}{(r-\alpha)A_0(\mathcal{E} + \lambda\tilde{\mathcal{H}})} \vee 1\right], \tag{10}$$

where all terms are as defined above. The formula clearly indicates that the value of the ecological benefit plus the health benefit hastens the time for policy adoption. It also indicates that delaying effects of rising adoption costs.

The table entitled "Elements of Decision and Optimal Timing" illustrates the elements of decision and optimal timing for a range of controls or levels of reduction. The parameter values are taken from the earlier exhibit "A numerical illustration," where the case of a 50% reduction was treated specifically, and contrasted with the case of a 40% reduction. The overall results show, for example, that reductions between 20% and 40% lead to relatively early adoption, but reductions near 10% or those beyond 50% lead to later adoption. At low levels of reduction, the ecological and health benefit is relative small, while at high levels of reduction, the adoption cost are steep. Incidentally, for the purpose of this illustration, adoption costs increase (exponentially) from a base of \$12.25 billion at a rate of 3.5% per percent reduction. Regardless of how one views this illustration politically in terms of the projection of adoption costs, the problem of projecting adoption costs econometrically is always among the critical problems of benefit-cost analysis.

Figure 1 provides a graphical illustration of optimal timing for the 40% level of control. The optimal timing is shown to be 2.5 years, which time maximizes the present value of the flow of social costs over all times τ using Equation 5 and the formula of Equation 10.

[8] The present work looks only at one-time reductions of emissions. Sometimes a sequential policy is warranted by various incentives for gradual policy adoption. For example, the traditional treatment foregoes the possibility that new information, whether it be scientific, economic or biomedical in nature, will come to light in ways that significantly impact the best regulatory response. Add the prospect of learning about the future or resolving some uncertainty and the value of flexibility takes on additional practical significance.

Elements of Decision and Optimal Timing *A Numerical Illustration*						
Level of Reduction	K^u	$A_0\mathcal{E}$	$\lambda A_0\tilde{\mathcal{H}}$	NPV_0	NPV_6	Optimal Timing T^{O}
10%	17.4	27.8	3.7	18.9	1.00	15.5 yrs
20%	24.7	55.6	7.1	49.8	-0.13	2.0 yrs
30%	35.0	83.3	10	80.4	-0.59	0.0 yrs
40%	49.7	111	13	111	-0.10	2.5 yrs
50%	70.5	139	16	141	1.71	7.7 yrs
60%	100	167	19	171	5.40	14.5 yrs
70%	142	194	21	201	11.7	22.5 yrs
80%	201	222	23	230	21.9	31.3 yrs
90%	286	250	25	260	37.3	40.7 yrs
100%	405	278	27	289	60.3	50.6 yrs

These summary calculations are in accord with the context given in "A numerical illustration." For each level of reduction, the decisional elements are expressed in billions of dollars, while the optimal timing of policy, T^O, is expressed in years. Each combination of level of reduction and optimal timing constitutes a control policy. The situation depicted at the row labeled 50% reduction recalls the details behind the earlier numerical illustration, where the claim was made for optimal timing at year 7.7 in the face of the various costs and benefits. Notice that NPV_6 is negative whenever year 6 surpasses optimal timing. This indicates that the incentive to wait has passed. Figure 1 provides a graphical illustration behind the maximization leading to the control policy as associated with the 40% reduction at year 2.5. Figure 1 also illustrates the optimal policy, taken over all reductions and times, to be about the 74% reduction at year 26. The details for this policy lie roughly between the constellation of decisional elements in the rows labeled by the 70% and 80% level of reduction.

Often, however, policy makers might aim at simultaneous optimization of timing and scope. This is accomplished by maximizing the value function over u and τ. As depicted in Figure 1, the present scenario yields approximately 74% for the optimal level of reduction with a corresponding timing at year 26.

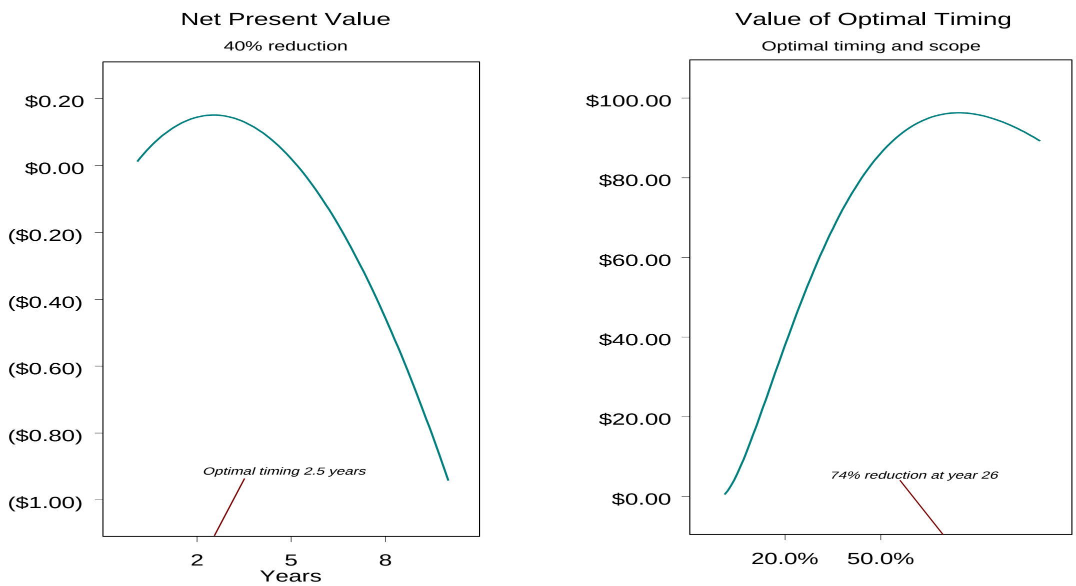

Figure 1. *The lefthand exhibit shows the graph of the net present value, $\tau \to NPV_\tau$, associated with the 40% level of reduction. Though NPV_0 is a positive* \$11 *billion, there is still value in waiting for the optimal time, T^O, at about 2.5 years. The righthand graph shows the value of an optimally timed policy over all levels of reduction; the value function itself is given in the next section as the present value of the flow of social costs. For a decision taken at the optimal time, the optimal control, u^*, is indicated at the 74% level of reduction. The corresponding optimal timing is year 26.*

Such is optimal policy in a world of certainty. The balance of this work treats the problem of optimal timing and scope of environmental health policy in the context of the simplified model of benefit-cost uncertainty.

Optimal policy in a world of uncertainty: Timing and scope of emissions reduction

This section describes the decision rule for optimal policy for the reduction of emissions of airborne particulate matter in a world of benefit-cost uncertainty. Policy is characterized by the optimal timing of an optimal reduction of emissions, as policy makers contemplate one of three actions: to act now, to wait until a later time or to never act at all. The timing of policy depends on the stochastic evolution of unit costs of emissions, while the scope of policy depends on maximizing the expected payoff at the time of adoption of policy.

Formally, each policy is determined by its timing and scope. This refers to a pair of numbers, namely (u, T), indicating a fractional reduction of emissions by an amount u as imposed at time T. The timing, T, denotes a stopping time that marks the beginning of a regulatory epoch as given by the interval $[T, \infty)$. The scope, u, of policy is a fractional reduction of emissions for the duration of the regulatory epoch. Recall that policy has the effect of both reducing emissions by $100u\%$ and of reducing the rate of loss-of-health production by $100(1 - \mathrm{e}^{-du})\%$. These effects lower the equilibrium level of the stock of pollution to $(1-u)\epsilon/\delta$, and the equilibrium level of loss-of-health stock to $\tilde{\eta}\mathrm{e}^{-du}/\tilde{\gamma}$.

Notice that, as the control variable in the problem, the emissions rate of particulates is the direct object of regulatory action. Health effects may be considered knock-on effects of the reduced emissions. Readers familiar with the National Research Council's Report to Congress, NRC 1998, recognize that the magnitude or even the presence of health effects constitutes a key source of scientific uncertainty for airborne particulate matter. For this reason, decision-makers may choose to investigate the effects on policy of a range of health effects in a sensitivity analysis.[9]

The next proposition describes the decision rule for the optimal timing of policy with an a priori fixed scope in a world of economic uncertainty.

11 Proposition: Optimal policy in a world of uncertainty. *Consider the case of a fixed level of reduction, u, for $0 < u \leq 1$. Suppose that economic uncertainty, σ^a, lies within $\sqrt{2\alpha}$, the square root of twice the appreciation rate. Let K^u denote the strictly positive expected cost of adoption and $p, \hat{\mathcal{E}}$, and $\hat{\tilde{\mathcal{H}}}$ the following set of coefficients:*

$$p = \frac{-(\alpha - (\sigma^a)^2/2) + \sqrt{(\alpha - (\sigma^a)^2/2)^2 + 2r(\sigma^a)^2}}{(\sigma^a)^2} > 1;$$

$$\hat{\mathcal{E}} = \frac{p-1}{p}\frac{u\epsilon}{(r-\alpha)(r+\delta-\alpha)} \qquad \textit{and} \qquad \hat{\tilde{\mathcal{H}}} = \frac{p-1}{p}\frac{(\tilde{\eta} - \tilde{\eta}e^{-du})}{(r-\alpha)(r+\tilde{\gamma}-\alpha)}.$$

Notice that p depends explicitly on benefit-cost uncertainty, namely σ^a, and that $\hat{\mathcal{E}}$ and $\hat{\tilde{\mathcal{H}}}$ represent simply rescaled versions of the decisional elements $\mathcal{E}$ and $\tilde{\mathcal{H}}$ of Equation 8.

[9] This is done in part below. Another approach, beyond the scope of this work, makes health effects the object of a (Bayesian) learning policy.

To describe policy, let $\mathcal{D}$ denote the open interval given by $(0, K^u/(\hat{\mathcal{E}} + \lambda\hat{\tilde{\mathcal{H}}}))$. The set $\mathcal{D}$ is the continuation region for the problem. The optimal timing of a level-u reduction of emissions is given by the stopping time $T_{\mathcal{D}}$, where

$$T_{\mathcal{D}} = \inf_{t>0}\{A_t \notin \mathcal{D}\} = \inf_{t>0}\{A_t \geq K^u/(\hat{\mathcal{E}} + \lambda\hat{\tilde{\mathcal{H}}})\},$$

namely the first time that the unit-cost of emissions, A, attains or exceeds the trigger cost as given by the cost-benefit ratio $K^u/(\hat{\mathcal{E}} + \lambda\hat{\tilde{\mathcal{H}}})$.

Finally, the value function for the problem, ϑ, is given by the function:

$$\begin{aligned} \vartheta(a, \tilde{h}, m) \quad = \quad & \frac{K^u}{p-1}\left(\frac{a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}}}{K^u}\right)^p - \frac{am}{r+\delta-\alpha} - \frac{a\epsilon}{(r-\alpha)(r+\delta-\alpha)} \\ & - \frac{a\lambda\tilde{\eta}}{(r-\alpha)(r+\tilde{\gamma}-\alpha)} - \frac{a\lambda\tilde{h}}{r+\tilde{\gamma}-\alpha}, \end{aligned}$$

for $m > 0$, $\tilde{h} > 0$ and a in $\bar{\mathcal{D}}$, the closure of $\mathcal{D}$. For a in $I\!R_+ \setminus \bar{\mathcal{D}}$, ϑ equals the payoff function g, where g is given by

$$g(a, \tilde{h}, m) = -\frac{a(1-u)\epsilon}{(r-\alpha)(r+\delta-\alpha)} - \frac{am}{r+\delta-\alpha} - \frac{a\lambda\tilde{\eta}e^{-du}}{(r-\alpha)(r+\tilde{\gamma}-\alpha)} - \frac{a\lambda\tilde{h}}{r+\tilde{\gamma}-\alpha} - K^u.$$

The payoff is simply the value of policy upon immediate adoption. □

Proposition 11 describes the policy for optimal timing of emissions reduction having a priori fixed scope. Policy makers monitor unit costs of particulate emissions until such time as they appreciate to the level of the trigger cost, $K^u/(\hat{\mathcal{E}} + \lambda\hat{\tilde{\mathcal{H}}})$, or the critical cost-benefit ratio. Policy is thus triggered when the marginal opportunity benefits of action offset the marginal opportunity costs of delay. The actual rate of unit-cost appreciation is uncertain, so the trigger cost involves an inflation factor, $p/(p-1)$. The inflation factor is particularly sensitive to the value of the parameter σ^a or the amount of benefit-cost uncertainty. The parameter p tends to 1 from above as uncertainty increases, thus inflating the trigger cost needed for adoption of policy. This in turn pushes the timing of policy (statistically speaking) further in the future. The origin of the parameter p emerges below, as does the origin of the leading term in value function, ϑ, for the problem. This latter term values society's option to adopt the policy based on current unit costs of emissions, a central notion in all irreversible decisions.

This solution is consistent with related results from the basic theory in Pindyck (1996). There is a new element, however, represented by the health-benefit factor, $\hat{\tilde{\mathcal{H}}}$, as the explicit accounting for the knock-on health effects of policy. A formal demonstration of the main result is given below, in part to draw attention to a theory of Brekke and Øksendal (1991, 1994) for optimal stopping and sequential stochastic control. This theory is of general interest to benefit-cost analysts because it is both broadly applicable and adapted especially to the needs of economic decision-making. The formal demonstration is followed in the next section by discussion of a numerical illustration of the main result. The illustration includes an investigation of the distribution of the optimal timing of policy. These results are of special interest to the quantitative policy analyst.

The first task is to show that the timing of policy falls within the class of optimal stopping problems considered for economic analysis in the work of Brekke and Øksendal (1991). To do so, let Y' denote the post-regulatory process, namely $(A, \tilde{H}, M)$ with constant emissions rate given by $(1-u)\epsilon$ and constant rate of loss-of-health production given by $\tilde{\eta}\mathrm{e}^{-du}$. These are, of course, the rates prevailing during the regulatory epoch. Next, let g denote the function on $\mathbb{R}^3_+$ given by

$$g(y) = \mathbb{E}'^y \int_0^\infty \kappa(Y'_t)\mathrm{e}^{-rt}dt - K^u, \tag{12}$$

where $\mathbb{E}'^y$ denotes conditional expectation with respect to the distribution of Y'. The entry point, y, is some triple $(a, \tilde{h}, m)$ in $\mathbb{R}^3_+$ and the function κ denotes the linear social costs of pollution as defined at Equation 4. The second term, K^u, as before denotes the expected cost of policy adoption as some function of u only. Now let Y denote the pre-regulatory processes, namely $(A, \tilde{H}, M)$ with constant emissions rate given by ϵ and constant rate of loss-of-health production given by $\tilde{\eta}$. These are, of course, the rates prevailing during the pre-regulatory epoch. For each policy (u, T), by virtue of the construction of the stock variables M and $\tilde{H}$, an elementary calculation reveals that the value function, $V^{u,T}$, of Equation 5 is given by the equation

$$V^{u,T}(y) = \mathbb{E}^y \left[\int_0^T \kappa(Y_t)\mathrm{e}^{-rt}dt + g(Y_T)\mathrm{e}^{-rT} \right], \tag{13}$$

where $\mathbb{E}^y$ denotes expectation with respect to the distribution of Y with entry point y, namely some entry point $(a, \tilde{h}, m)$ in $\mathbb{R}^3_+$ for the pre-regulatory process $(A, \tilde{H}, M)$.

Optimal timing is therefore equivalent to the optimal stopping problem for the pre-regulatory process Y with cost function κ, payoff g and discounting at rate r. For fixed reduction u, the value function in this case, $V^*(u)$, is given by the equation

$$V^*(u;y) = \sup_T \mathbb{E}^y \left[\int_0^T \kappa(Y_t) e^{-rt} dt + g(Y_T) e^{-rT} \right], \tag{14}$$

for every starting point y in $\mathbb{R}^3_+$. Proposition 11 exhibits the unique solution to this problem.

Demonstration of Proposition 11. The demonstration consists of verifying the conditions of Theorem 2 in Brekke and Øksendal (1991). Regarding the correspondence of notation, the payoff function g has the same designation in both settings. Otherwise, the functions h and f in Brekke and Øksendal correspond respectively to the functions ϑ and κ in this setting. The open sets W and $\mathcal{D}$ in Brekke and Øksendal correspond respectively to the open sets $(0,\infty)^3$ and $\mathcal{D} \times (0,\infty)^2$ in the present model. The fact that $-\kappa$ is always positive, and the monotone convergence theorem, replaces Brekke and Øksendal's Condition (60) on appeal to their Generalized Dynkin formula. For purposes below, let $\mathcal{A}$ denote the Markov Generator belonging to the process Y. Let L denote the Markov Generator $\mathcal{A} - r$ for the process Y with killing at rate r. The latter process corresponds to the process X in Brekke and Øksendal.

As suggested above, define W as the strictly positive octant of $\mathbb{R}^3$ and redefine the continuation region, $\mathcal{D}$, of the present model as a subset of this W. Notice first that the boundary, $\partial\mathcal{D}$, namely the set $\{(a, \tilde{h}, ma) : a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}} = K^u\}$, is manifestly C^1. At the same time notice that the Markov Generator $\mathcal{A}$ is elliptic on W. Therefore, on every sufficiently small ball centered on the boundary set $\partial\mathcal{D} \cap W$, the Markov Generator $\mathcal{A}$ is uniformly elliptic. Finally, starting from any point in W, the stopping time $T_{\mathcal{D}}$ is almost surely finite, since A is a Geometric Brownian motion and 2α strictly exceeds $(\sigma^a)^2$. This is demonstrated in detail for example in Karlin and Taylor (1975), page 362, for the Brownian motion lnA having strictly positive drift, given here by $\alpha - \frac{1}{2}(\sigma^a)^2$.

Next, as a polynomial on the closure of the continuation region, the candidate ϑ for the value function belongs to the set of smooth functions $C^1(\bar{\mathcal{D}}) \cap C^2(\mathcal{D})$, again with $\mathcal{D}$ viewed as a subset of $\mathbb{R}^3_+$. And as the payoff function g satisfies the equation

$$g(m, \tilde{h}, a) = -\frac{a(1-u)\epsilon}{(r-\alpha)(r+\delta-\alpha)} - \frac{am}{r+\delta-\alpha} - \frac{a\lambda\tilde{\eta}e^{-du}}{(r-\alpha)(r+\tilde{\gamma}-\alpha)} - \frac{a\lambda\tilde{h}}{r+\tilde{\gamma}-\alpha} - K^u,$$

for every $(a, \tilde{h}, m)$ in W, it too is evidently C^1 on W and C^2 on $W \setminus \bar{\mathcal{D}}$. At the same time, by definition of $\mathcal{D}$, a brief algebraic exercise shows that the function $\vartheta - g$ is positive on the open set $\mathcal{D}$.

Regarding the payoff function, it remains only to show that the function Lg is bound by $-\kappa$ on $W \setminus \bar{\mathcal{D}}$, where κ refers to the linear social cost of Equation 4 and where L is the Markov Generator $\mathcal{A} - r$. To wit, by definition of g and L, the function Lg is seen to satisfy the equation

$$Lg(m, \tilde{h}, a) = -\kappa(m, \tilde{h}, a) + rK^u - \frac{au\epsilon}{r+\delta-\alpha} - \frac{a\lambda(\tilde{\eta} - \tilde{\eta}e^{-du})}{r+\tilde{\gamma}-\alpha}.$$

Now apply the definition of $\mathcal{D}$ and use that $r > r - \alpha > 0$ to show algebraically that $Lg \leq -\kappa$ for every $(a, \tilde{h}, m)$ in $W \setminus \bar{\mathcal{D}}$.

It remains, finally, to demonstrate the optimal stopping criteria for the underlying free boundary problem, including the 'smooth pasting' and 'high contact' conditions. For this purpose, let $\mathcal{OV}$ denote the mapping given by $a \to (K^u/(p-1))\left((a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}})/K^u\right)^p$. Now apply the Markov generator L to ϑ and observe that the resulting function $L\vartheta$ satisfies the equation

$$L\vartheta(a, \tilde{h}, m) = [p(p-1)(\sigma^a)^2/2 + \alpha p - r]\mathcal{OV}(a) + am + a\lambda\tilde{h} = -\kappa(m, \tilde{h}, a),$$

which equation uses that $r = p(p-1)(\sigma^a)^2/2 + \alpha p$ and the definition of κ. Next, for $(a, \tilde{h}, m)$ in $\partial\mathcal{D}$, the function $\vartheta - g$ satisfies the equation

$$(\vartheta - g)(a, \tilde{h}, m) = \frac{K^u}{p-1} - \frac{p}{p-1}\left[a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}}\right] + K^u = 0,$$

since a equals $K^u/(\hat{\mathcal{E}} + \lambda\hat{\tilde{\mathcal{H}}})$ on the boundary. This verifies 'smooth pasting' or the continuity between the value function and the payoff function on the boundary of the continuation region. Finally, equating the partial derivatives $\partial_a\vartheta$ and $\partial_a g$, observe that

$$\frac{pK^u}{p-1}\left(\frac{a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}}}{K^u}\right)^p = \frac{p}{p-1}(a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}}),$$

for every point on the boundary $\partial\mathcal{D}$. This verifies the 'high contact' condition and completes the demonstration. $\square$

Such is policy for a priori fixed scope. Policy makers, on the other hand, may also be interested in the optimizing the scope as well the timing. The value function for optimal timing and scope, V^*, is given by the formula

$$V^*(a,\tilde{h},m) = \sup_{u,T} I\!E^{a,\tilde{h},m}\left[\int_0^T X^{u,T}\mathrm{e}^{-rt}dt + g(A_T,\tilde{H}_T,M_T;u)\mathrm{e}^{-rT}\right], \tag{15}$$

for every starting point $(a,\tilde{h},m)$ in $I\!R^3_+$, where the payoff function, g, is shown to be explicitly dependent on the scope, u, of policy. The solution to this problem may be described as a modification of the solution to the optimal stopping problem in Proposition 11.

That is, for each a, m, and $\tilde{h}$, notice from Proposition 11 that the function $u \to g(m,\tilde{h},a;u)$ is a continuous function over the compact domain $[0,1]$. There is therefore a point $\bar{u}(a)$ in $[0,1]$ such that $g(m,\tilde{h},a;\bar{u}(a))$ satisfies the equation

$$g(a,\tilde{h},m;\bar{u}(a)) = \max_u g(a,\tilde{h},m;u).$$

The mapping $a \to \bar{u}(a)$ then records the payoff maximizing scope, which can at most depend on a and the ecological and health benefit factors $\mathcal{E}$ and $\tilde{\mathcal{H}}$. The function $\bar{u}$ may be shown to be continuous and monotone, but it may possess points of non-differentiability. This may happen, for example, whenever 0% or 100% reductions are optimal over a range of costs per unit emissions. These points of non-differentiability in turn compromise the (everywhere) differentiability of the mapping $\bar{g}$, namely $(a,\tilde{h},m) \to g(a,\tilde{h},m;\bar{u}(a))$, which is the payoff attained at optimal scope. The same issue regarding differentiability arises for the value function ϑ of Proposition 11, once replaced by the new value function $\bar{\vartheta}$, namely the value attained by simultaneous optimization of timing and scope.

There is a new payoff-maximizing trigger cost to be constructed by first replacing the parameter u with the function $\bar{u}$ in the coefficients $\hat{\mathcal{E}}$, $\hat{\tilde{\mathcal{H}}}$ and K^u. This way the trigger cost is given by the benefit-cost ratio taken at optimal scope. The claim is that the optimal timing of policy is given by the first time that the unit-cost, A, of emissions reaches a trigger cost, a^* say, at which time the optimal scope of policy, u^* say, is given by the optimal level of reduction $\bar{u}(a^*)$.

A formal demonstration of this result in not given here. The natural approach to a formal demonstration, however, is to pass first to the equivalent stochastic control problem over a class of Markov controls. The demonstration would then follow the outline of the statement and proof of the converse to the celebrated Hamilton-Jacobi-Bellman equation.[10]

Optimal policy in a world of uncertainty

This section illustrates numerically the basic elements of decision-making for the optimal timing and scope of policy under uncertainty. The values for the elementary parameters are given by those of earlier numerical illustrations: $r, \alpha, \tilde{\gamma}, \delta, \epsilon, \lambda, d,$ and η are given by 0.04, 0.025, 0.025, 0.045, 10 megatons per year, 200,000 unit cost loss-of-health per unit cost emissions, 7% health effect per percent reduction, and 6.5 units loss-of-health per year, respectively. Adoption costs appreciate exponentially from a base of \$12.25 billion at a rate of 3.5% per percent reduction. The level of uncertainty, σ^a, is typically given by 20%, but it will be varied below. So will the elasticity, d, for the health effect.

The most basic decisional element is the continuation region for optimal timing of policy. For an a priori fixed level of reduction, the continuation region is given by an interval between \$0.00 and $\$pK^u/(p-1)(\mathcal{E}+\lambda\tilde{\mathcal{H}}))$ per ton of emissions per year. For purposes of decision making, policy analysts would monitor unit cost of emissions, namely the process $t \to A_t$, until such time as it first reaches or exceeds $\$pK^u/(p-1)(\mathcal{E}+\lambda\tilde{\mathcal{H}}))$, the so-called trigger cost for the adoption of policy. As derived from the table "Elements of Decision and Optimal Timing," for example, the trigger cost for a 50% reduction is given by the figure \$49.95 per ton per year and for a 40% reduction by the figure \$43.84 per ton per year. The inflation factor, $p/(p-1)$, is 4.39.[11]

[10] One caveat is that a formal demonstration would involve a venture into the technical territory of a generalized Dynkin formula for stochastically $\mathcal{C}^2$ functions, because the value and payoff functions may not be differentiable everywhere. The required theory may be found, for example, in the work of Brekke and Øksendal (1994). In this theory the differentiability of the value and payoff functions need only hold almost everywhere with respect to the Green's measure for the problem. The exceptional sets in the present case refer simply to sets of Lebesgue measure zero on the positive half line.

[11] Recall that optimal timing in a world of certainty leads to a deterministic stopping time given by that time τ such that $A_0 e^{\alpha\tau}$ first reaches or exceeds $\$rK^u/(r-\alpha)(\mathcal{E}+\lambda\tilde{\mathcal{H}})$ per ton per year. Thus, the trigger

The first exit time from the continuation region is, of course, a random variable that depends on the stochastic behavior of A. For the Geometric Brownian motion it is possible to calculate analytically the distribution of the optimal timing of policy. Figure 2 illustrates two such distributions for the timing a 40% reduction, one for 11% uncertainty and the other for 16% uncertainty. These graphs depict densities of the so-called inverse Gaussian distribution, where this application exhibits two members of a family of distributions marked by a strong dependence of both the mean and the variance on the single parameter, σ^a, for economic uncertainty. As σ^a decreases, the distribution of optimal timing of policy concentrates peakedly about 2 to 3 years. As σ^a increases toward $\sqrt{2\alpha}$, in contrast, the distribution becomes increasingly heavy tailed while the mean tends toward infinity.

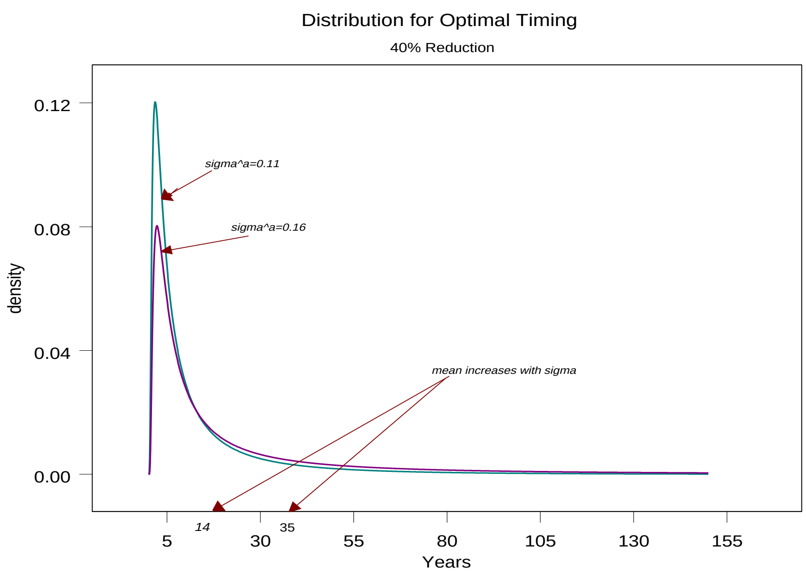

Figure 2. *This graph depicts the density of the distribution of the optimal timing of a 40% reduction for σ^a equal 11% and 16%. These distributions exhibit instances of what are commonly called inverse Gaussian distributions. These distributions illustrate a family of cases in which the mean and variance are functionally related through the parameter σ^a. As σ^a decreases, the distribution becomes more peaked near 2 or 3 years, the timing of policy in a world of certainty from the table "Elements of Decision and*

cost for a 40% reduction in a world of certainty is about \$26 per ton per year. This is about 40% less than the \$43.84 trigger cost in a world of 20% uncertainty.

Optimal Timing." As σ^a increases, the distribution grows increasingly heavy tailed and the expectation becomes unbounded.

The distribution for optimal timing as well as expressions for the mean and variance are easily derived using martingale methods. This is illustrated for example in Karlin and Taylor (1975), pages 357-363. For $\sigma^a < \sqrt{2\alpha}$ and for $pK^u/(p-1)(\mathcal{E} + \lambda\tilde{\mathcal{H}}) \geq A_0$, the expected timing of a level-u reduction is given by

$$\mathbb{E}_0 T^o = \ln\{pK^u/(p-1)(A_0\mathcal{E} + A_0\lambda\tilde{\mathcal{H}})\}/(\alpha - (\sigma^a)^2/2), \tag{16}$$

while the variance of timing is given by

$$\mathrm{Var} T^o = (\sigma^a)^2 \ln\{pK^u/(p-1)(A_0\mathcal{E} + A_0\lambda\tilde{\mathcal{H}})\}/(\alpha - (\sigma^a)^2/2)^3. \tag{17}$$

These expression show, for example, that the expected timing of policy rises to ∞ as σ^a approaches the critical point $\sqrt{2\alpha}$. Although the mean and the variance provide some expectation for the optimal timing of policy, notice from Figure 2 that the distribution of optimal timing is rather right skewed. this suggests that the median and interquartile range would provide a better indication of typical timing and its variability. The quartiles, however, need to be approximated. An analytical approximation, suggested in Johnson, Kotz and Balakrishnan (1994), pages 261-270, is given by

$$q(v; T^o) = \mathbb{E}_0 T^o \mathrm{e}^{(\Phi^{-1}(v) - 0.5\sqrt{(\sigma^a)^2/(\alpha - (\sigma^a)^2/2)})\sqrt{(\sigma^a)^2/(\alpha - (\sigma^a)^2/2)})} \tag{18}$$

where Φ^{-1} denotes the inverse of the standard Gaussian distribution function. Simply choose v equal to 0.5, 0.25, and 0.75 for the median, the lower quartile and the upper quartile respectively.

Some discretion is required in applying this approximation, because it is particularly sensitive to the magnitude of the coefficient of variation. For present purpose, the approximation is satisfactory for σ^a within about 16%. Beyond that point, it is better to rely a numerical approximation using the underlying distribution function. Figure 3 illustrates the result of applying the approximation at Equation 18 for a 40% reduction. As σ^a ranges from 0.5 to 16% uncertainty, the median timing of policy increases from a low of about 2 to 3 years, roughly the timing in a world of certainty, to a high of about 12 years or so. Twelve years and four months corresponds roughly to the median timing at 22.4% uncertainty. This level of uncertainty corresponds to the critical case, $\sigma^a = \sqrt{2\alpha}$, at which the mean timing is infinite. For modest uncertainty, say σ^a in the range between 5 and 8%, notice that the middle half of the distribution of optimal timing for a 40% reduction lies more or less in the interval between about 5 and 11 years. As the level of economic uncertainty increases from 8% to 16%, in contrast, the interquartile range increases some two-to-six fold.

The results of Figure 3 illustrate an example of a sensitivity analysis for optimal timing as a function of uncertainty. Policy makers may also use this device for other elementary parameters and show, for example, that the median timing of policy is a decreasing function of the health benefit. An example of such an effect is examined below in the context of simultaneous optimization of the timing and scope of policy.

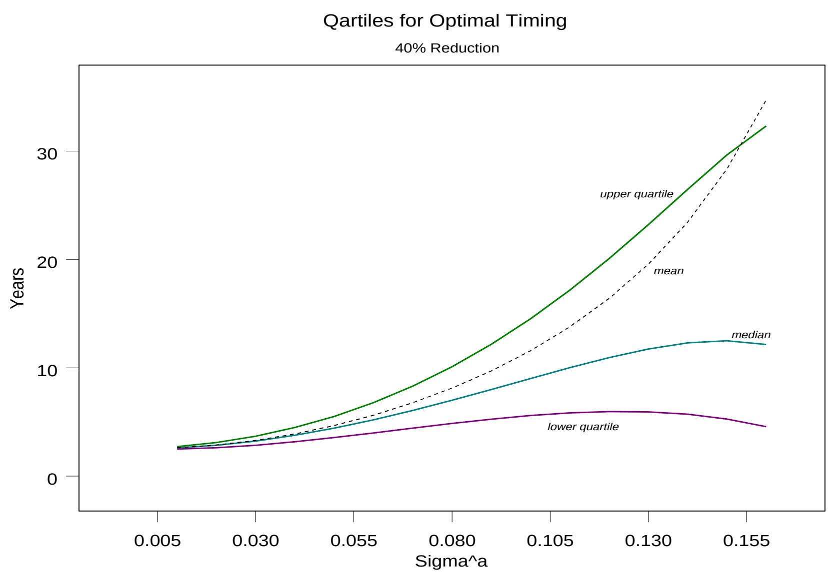

Figure 3. *This figure graphs the quartiles of the distribution of optimal timing for a 40% reduction as a function of uncertainty. The quartiles are based on the approximation given by Equation 18. Uncertainty ranges from a low of about 0.5% to a high of about 16%, beyond which point the approximation of Equation 18 breaks down. As uncertainty increases, the median timing rises from a low of about 2 to 3 years to a high of about 12 years. Twelve years corresponds roughly to the median timing at 22.4% uncertainty or the point at which the mean timing is infinite. For uncertainty within about 8% or*

so, the middle half of the distribution lies within a range of no more than about 6 to 8 years. For greater uncertainty, the spread of the distribution increases considerably to as much as 20 to 30 years. The graph of the mean timing of policy is also shown for reference. The mean, of course, increases rapidly to ∞ *as* σ^a *approaches* $\sqrt{2\alpha}$ *or 22.4%.*

The extremal case, where $\sigma^a = \sqrt{2\alpha}$, is a special one in which the expected timing is actually infinite; timing itself, though, is finite with probability one. The median timing of policy, on the other hand, is given by

$$\text{(19)} \qquad \mathrm{Med}T^o = (\ln\{pK^u/(p-1)(A_0\mathcal{E} + A_0\lambda\tilde{\mathcal{H}})\})^2/(\sigma^a)^2\Phi^{-1}(3/4).$$

The lower quartile is given by replacing the fraction 3/4 in Equation 19 with 7/8, the upper quartile by replacing 3/4 with 5/8. These expressions may be derived from the reflection principle for the Brownian motion; see for example Karlin and Taylor (1975), pages 346-347. For $\sigma^a > \sqrt{2\alpha}$ the optimal timing may technically be infinite with positive probability. In this event, the demonstration of Proposition 11 breaks down.

The value and payoff functions are the next basic elements of decision-making under uncertainty. Figure 4 depicts the graphs of the value function, ϑ, and the payoff function, g, of Proposition 11 for an emissions reduction policy of 50%. Notice that the value function dominates the payoff function for cost below \$54.30 per ton per year, the point of high contact and smooth pasting. This is where the benefits of emissions reduction to the ecology and human health combine to equal the social costs of adoption. It is also where the marginal payoff of policy exactly offsets the marginal value of holding the option to adopt policy. In the theory of Proposition 11, though not depicted in Figure 4, the value function technically flattens out and follows along the graph of the payoff function for costs in excess of the trigger cost.

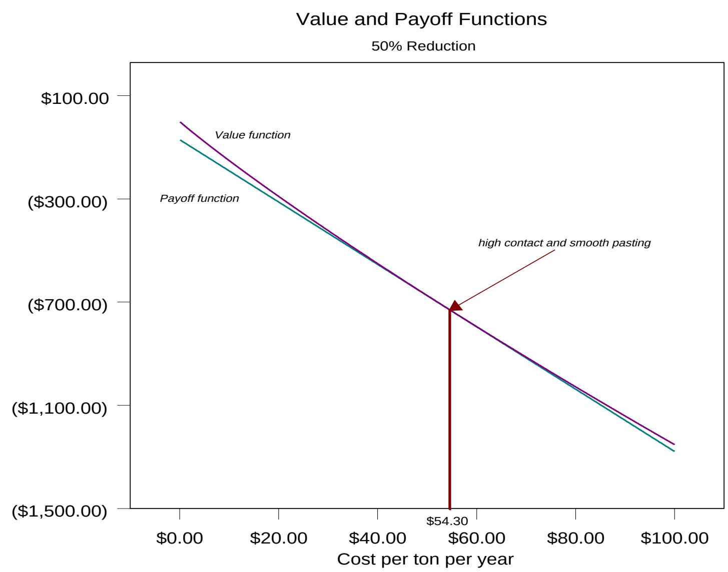

Figure 4. *This figure graphs the value function and the payoff function for decision-making under uncertainty as a function of the cost of emissions per ton per year. The level of uncertainty, σ^a, is given by 20%. The initial values for the pollution stock, m, and the loss-of-health stock, $\tilde{h}$, are given by their respective equilibrium values of 222 megatons and 260 units loss-of-health. The value function dominates the payoff function for costs in the continuation interval, namely* $[0, \$54.30)$. *Policy is triggered when costs first reach or exceed \$54.30 per ton per year, the point of high contact and smooth pasting. Thereafter, technically, though not depicted graphically, the value function tracks the payoff function. The nonlinearity in the value function is a measure of the value of society's option to adopt the emissions reduction policy, that is, the so-called option value.*

Option value refers to a measure of the value of society's option to reduced emissions at such time as costs warrant the adoption of policy. This measure refers analytically to the function, $\mathcal{OV}$ say, given by the equation

$$\mathcal{OV}(a) = \frac{K^u}{p-1}\left(\frac{a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}}}{K^u}\right)^p, \tag{20}$$

for all costs a per ton of emissions per year. The influence of option value is evidenced in Figure 4 by the curvature in the value function. Its principle determinant is the characteristic root, p, and the benefit-cost ratio, namely $(a\hat{\mathcal{E}} + a\lambda\hat{\tilde{\mathcal{H}}})/K^u$, of the sunk benefits (ecological plus health) of policy to the sunk cost of adoption. he irreversibilities behind these sunk benefits and cost were first introduced in a world of certainty, where a similar ratio, without the effects of uncertainty, played the important role in optimal timing of emissions reduction.

Notice that the motivation for optimal timing of policy is present in both worlds of certainty and uncertainty and, in essence, for the same reasons: the presence of irreversibility. The presence of economic uncertainty, for its part, serves to make the timing of policy uncertain and to inflate the ultimate cost needed to trigger its adoption. The latter effect is derived in part from the option value with its introduction of the parameter p and, more importantly, the inflation factor, $p/(p-1)$, itself an increasing function of economic uncertainty.

Finally, consider the simultaneous optimization of timing and scope of policy and the payoff-maximizing reduction. The payoff-maximizing reduction refers analytically to the function $\bar{u}$ given by the equation

$$\bar{u}(a) = \arg\max_u g(m, \tilde{h}, a; u), \tag{21}$$

where g is the payoff function of Proposition 11. While $\bar{u}$ will often be a continuous function, it will generally have points of non-differentiability, accounting partly for the need of a more complicated mathematical treatment of policies for optimal timing and scope.

Figure 5 graphs numerical approximations to two payoff-maximizing reductions as a function of the cost of emissions. In the top panel, the level of uncertainty is 20% and the health effect parameter is 0.7 as before. In the bottom panel, in contrast, the level of uncertainty is 10% and the health effect parameter is 7 or an order of magnitude larger than before. The payoff maximizing reduction rises fairly steeply in each case to 100% somewhere between \$125 and \$130 per ton per year. Beyond these costs, the payoff maximizing reduction remains 100%. In scenario one, characterized by high uncertainty and a small (relatively) health effect, optimal policy triggers at about \$158.70 per ton per year and the optimal scope is 100%. If initial costs of emission per ton per year are about \$25 per ton per year, the median timing of policy approaches 126 years. In scenario two, characterized in contrast by half the uncertainty and ten times the health effect, optimal policy triggers at about \$46.80 per ton per year and optimal scope is about 72%. If initial costs of emissions are about \$25, the median timing of policy in scenario two is closer to 24 years.[12]

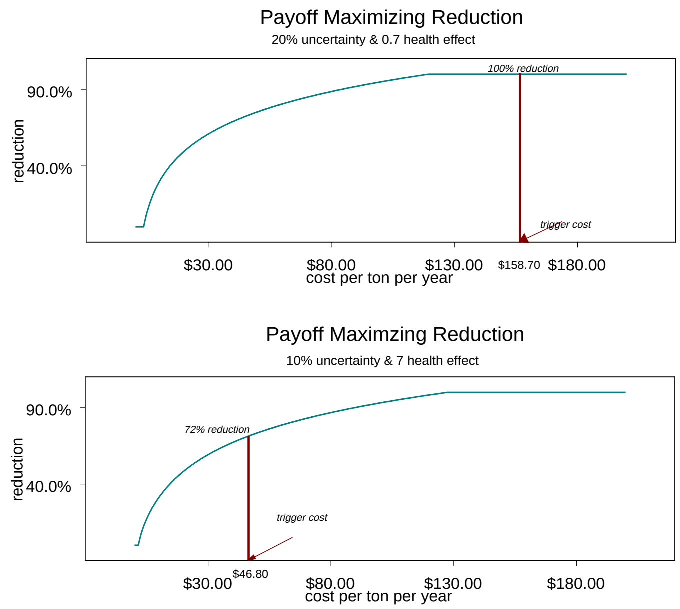

Figure 5. *This figure graphs the payoff maximizing reduction, $\bar{u}$, under two scenarios. The top panel depicts the case where uncertainty is 20% and the health effects parameter is set to its normative value 0.7. The bottom panel depicts the case where uncertainty is 10% and the health effects parameter is 7, that is half the uncertainty and ten times the health effect of scenario one. Each curve depicts a smooth, increasing function of costs with at least one point of non-differentiability. That cost is roughly \$120 per ton per year in scenario one and roughly \$130 per ton per year in scenario two. Optimal policy under scenario one triggers at \$158.70 per ton per year with 100% the optimal scope. Optimal policy under scenario two, in contrast, triggers at \$46.80 with 72% the optimal scope. If initial costs are about \$25 per ton per year, the median time to policy adoption in scenario one is 126 years. The median timing in scenario two is in contrast 24 years. The present value of adoptions costs under the two scenarios would then be about \$36 billion and \$60 billion, respectively.*

[12]The policy differences between scenario one and two are largely due to the difference in uncertainty. That is, for 20% economic uncertainty, large increases in the health effect bring about only modest changes in timing and no change in scope. But the combined effect of reducing uncertainty and, say, discovering larger health effects of policy is considerable.

In each of these two scenarios the timing of policy is motivated in part by the desirability of lowering the present value of the adoption costs. These costs are approximately \$400 billion for the 100% reduction and \$152 billion for the 72% reduction. Using the Laplace transform of the optimal timing of policy from Equation (5.5), page 362, in Karlin and Taylor (1975), the expected present value of the adoption costs are given by \$36 billion and \$60 billion, respectively. Though larger in scenario two, the tradeoff revolves around the accumulated years of sunk benefits to the ecology and human health brought about (statistically speaking) by earlier adoption of policy. Nevertheless, either policy may best be seen as a long term goal rather than a one-time rule. Given the many uncertainties here, prudence might prefer to resolve some uncertainty with for example better estimates of health effects. Perhaps a sequential policy would be a course worth investigating beyond the scope of the present work.

Discussion and continuing work

This work presents a simplified model for timing and scope of emissions reductions for airborne particulate matter under conditions of economic uncertainty. The model is derived from the basic theory for the economic analysis of environmental policy developed by Pindyck (1996). An added feature of the present model is the representation of health effects, which were added to ecological effects as an additional source of sunk benefits of regulation. A numerical illustration of the main results was given not for the purpose of prescribing policy, but for the purpose of introducing quantitative policy analysts to the nature, potential and scope of Pindyck's basic theory for the analysis of environmental health regulation.

The model is naturally an idealization of the problem. It does allow abstraction of the essential elements of decision-making under uncertainty, best used for establishing guidelines for policy makers, rather than for setting binding rules. One effective use of these results is to subject each of the decisional elements to a sensitivity analysis over the relevant parameters– such as σ^a and p for $\mathcal{OV}$, and so on. Another effective use is in the manner suggested by Cox (1997), where it serves as a benchmark against which to refer the results of numerical simulations of more detailed models.

One direction for continuing research is to introduce stock effects of the pollutant, the loss-of-health or both on optimal policy. Expressions of stock effects significantly complicate the analytical treatment of the problem, but some expressions may be tractable to numerical solution. A second direction for continuing research concerns the representation of health effects, specifically in light of modern developments in the econometric valuation of health and in light of the modern approaches to the epidemiology of upper respiratory disease.

Health effects may be chief among the three sources of uncertainty—economic, ecological and biomedical—in the analysis of regulation for airborne particulate matter. In recommending the $PM_{2.5}$-rule, the EPA was acting on recent epidemiological evidence linking particulate matter to adverse health effects, particularly asthma and other upper respiratory diseases. Setting the new standards required that EPA officials forge a policy decision with potentially far reaching implications for the U.S. economy from their assessment of the best available scientific evidence. "Even as the new standards were being promulgated," the National Research Council's Committee on Research Priorities for Airborne Particulate Matter reported in 1998, "scientists and policymakers recognized that further research was needed to address key uncertainties." In addition to scientific uncertainty about the spatio-temporal characteristics of airborne pollution, a critical, medical uncertainty remained about the exact health effects if any of human exposure to small-sized pollution particles.

This suggests a third, perhaps more vital, direction for continuing research. There is a need in the present case for sequential approaches or learning policies for the regulation of airborne particulate matter. This way significant uncertainties, such as the magnitude and scope of health effects, can be addressed gradually and scientifically as part of an overall policy response to the problem.

Acknowledgments. Warm hospitality and generous support for this research is acknowledged from the *National Research Center for Statistics and the Environment*, University of Washington, Seattle. Support is also acknowledged from Chapman University's Office of Faculty Research and Development. I wish to thank Peter Guttorp especially for making the results of Greg Nothstein's Master's paper on PM regulation available to me.

References

[1] KJ Arrow, ML Cropper, GC Eads, RW Hahn, LB Lave, RG Noll, PR Portney, M Russell, R Schmalensee, VK Smith, and RN Stavins. Is there a role for benefit-cost analysis in environmental, health and safety regulation? *Science*, 272:221–222, 1996.

[2] KJ Arrow and AC Fisher. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics*, 88(2):312–19, May 1974.

[3] JR Birge and CH Rosa. Modeling investment uncertainty in the costs of global CO_2 emission policy. *European Journal of Operations Research*, 83:466–488, 1995.

[4] KA Brekke and B Oksendal. The high contact principle as a sufficiency condition for optimal stopping. In D Lund and B Oksendal, editors, *Stochastic Models and Option Values: Application to Resources, Environment and Investment Problems*, pages 187–208, Amsterdam, 1991. North Holland.

[5] KA Brekke and B Oksendal. Optimal switching in an economic activity under uncertainty. *Siam Journal of Control and Optimization*, 32:1021–1036, 1994.

[6] JM Conrad. Global warming: When to bite the bullet. *Land Economics*, 73(2):164–173, May 1997.

[7] National Research Council. *Research Priorities for Airborne Particulate Matter: I: Immediate Priorities and a Long-Range Research Portfolio.* National Academy Press, Washington DC, 1998.

[8] DR Cox. The current position of Statistics: A personal view. *International Statistical Review*, 65:261–290, 1997.

[9] ML Cropper. Health, investment in health and occupational choice. *Journal of Political Economy*, 85(6):1273–1294, 1977.

[10] ML Cropper. Measuring the benefits from reduced morbidity. *American Economic Review*, 71(2):235–240, 1981.

[11] AC Fisher and JV Krutilla. Economics of nature preservation. In AV Kneese and JL Sweeney, editors, *Handbook of Natural Resources and Energy Economics: Volume I*, pages 165–189, New York, NY, 1985. Elsevier Publishers B.V.

[12] AC Fisher, JV Krutilla, and CJ Cicchetti. The Economics of Environmental Preservation: A Theoretical and Empirical Analysis. *American Economic Review*, 62:605–19, 1972.

[13] M Grossman. On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255, 1972.

[14] WM Hanemann. Information and the concept of option value. *Journal of Environmental Economics and Management*, 16:23–37, 1989.

[15] C Henry. Investment decisions under uncertainty: The 'irreversibility' effect. *The American Economic Review*, 64(6):1006–12, December 1974.

[16] NL Johnson, S Kotz, and N Balakrishnan. *Continuos Univariate Distributions: Volume 1, Second Edition.* John Wiley & Sons, New York, NY, 1994.

[17] S Karlin and HM Taylor. *A First Course in Stochastic Processes.* Academic Press, New York, NY, 1975.

[18] WD Nordhaus. To slow or not to slow: The economics of the greenhouse effect. *The Economic Journal*, 101:920–937, 1991.

[19] MJ Phelan. Environmental health policy decisions: The role of uncertainty in economic analysis. *Journal of Environmental Health*, 61:8–12, 1998.

[20] RS Pindyck. Sunk costs and sunk benefits of environmental policy: I. Basic theory. Technical report, Massachusetts Institute of Technology, Cambridge Massachusetts, 02139, 1996.