

NBER WORKING PAPER SERIES

CONTROLLING STOCKS AND FLOWS TO PROMOTE QUALITY:
THE ENVIRONMENT, WITH APPLICATIONS TO
PHYSICAL AND HUMAN CAPITAL

Nathaniel O. Keohane
Benjamin Van Roy
Richard J. Zeckhauser

Working Paper 7727
<http://www.nber.org/papers/w7727>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2000

The authors thank, without implicating, George Borjas, David Ellwood, Eduardo Engel, Peter Glynn, Sandy Jencks, David Laibson, John Leahy, Pete Veinott, and participants in the Management Science and Engineering Colloquium at Stanford University. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by Nathaniel O. Keohane, Benjamin Van Roy, and Richard J. Zeckhauser. All rights reserved. Short sections of text not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Controlling Stocks and Flows to Promote Quality: The Environment,
With Applications to Physical and Human Capital
Nathaniel O. Keohane, Benjamin Van Roy, and Richard J. Zeckhauser
NBER Working Paper No. 7727
June 2000
JEL No. C61, E22, Q20

ABSTRACT

Our analysis melds two traditional approaches to promoting quality. The first is restoring the stock of quality. The second is curbing its flow of deterioration. Although both approaches are widely used in real world settings, analytic models have tended to focus on one strategy or the other.

We consider a class of problems, which we call “SFQ” problems, in which both *stocks* and *flows* can be controlled to promote *quality*. We develop our results in the context of environmental quality, drawing on real-world examples from atomic wastes to zebra mussels. But the lessons are general, and we show how they apply to promoting the quality of both physical and human capital.

We first study optimal policies in the limiting cases when only abatement or restoration is possible. We then focus on the full SFQ world, where both approaches can be used. We show that the optimal policy employs both instruments. Moreover, when combined optimally, neither strategy takes the form it would in the absence of the other.

Nathaniel O. Keohane
Graduate School of Arts and Sciences
Harvard University
Cambridge, MA 02138
nkeohane@fas.harvard.edu

Benjamin Van Roy
Management Science and Engineering
Stanford University
Mail Code 4026
Terman Engineering B
Stanford, CA 94305
bvr@stanford.edu

Richard J. Zeckhauser
JFK School of Government
Harvard University
79 John F. Kennedy Street
Cambridge, MA 02138
and NBER
richard_zeckhauser@harvard.edu

CONTROLLING STOCKS AND FLOWS TO
PROMOTE QUALITY:
THE ENVIRONMENT, WITH APPLICATIONS
TO PHYSICAL AND HUMAN CAPITAL

By Nathaniel O. Keohane, Benjamin Van Roy, and Richard J. Zeckhauser*

1 Introduction

This paper considers the problem of promoting quality in a dynamic setting. The quality of a valued resource diminishes over time; it may be restored

*

periodically, or the process of deterioration may be slowed. For example, environmental quality, which we represent as a stock, diminishes as solid waste accumulates at a landfill. The flow of waste may be slowed through recycling, composting, or waste reduction. Eventually, the landfill is capped and the quality of the site is restored. In the management of physical capital, quality diminishes as existing capital deteriorates and becomes obsolete. Investment in new capital restores the stock; maintenance slows depreciation. Human capital behaves in much the same fashion. An engineer whose skills have become obsolete can be “traded in” through dismissal or downsizing, and that “trade-in” can be delayed through maintenance efforts, such as continued job training.

We meld two traditional approaches to promoting quality: controlling stocks and controlling flows. The first entails restoring the quality of a resource. The second involves curbing its deterioration. Although both approaches are widely used in real world settings, analytic models tend to focus on one strategy or the other. For example, pollution-control models typically concentrate on balancing the marginal benefits and costs of flows of pollution. Stock-control measures do not enter these models, or are treated separately. In contrast, models dealing with physical capital focus on investment and the stock of capital. Maintenance tends to be ignored, or is treated as exogenous. And while “trade-in” and “maintenance” strategies for human capital are widely discussed, they are not considered together in optimizing models. In each of these settings, both *stocks* and *flows* can be controlled to promote *quality*. We refer to this class of problems as “SFQ” problems. To facilitate exposition we focus on environmental quality, but the lessons are general. Our discussion of applications returns briefly to promoting the quality of physical and human capital.

In the well-behaved world of economics textbooks, the marginal benefit and cost curves of reducing pollution depend only on current quality, and slope, respectively, down and up. A constant level of environmental quality is maintained where the marginal benefit of reducing pollution, adjusted for the discount rate and the decay rate of the stock, equals the marginal cost of abating it.¹ Once the resource reaches such a steady state, optimal abatement efforts just keep up with net new accumulation. If there is uncertainty about flow or decay rates, environmental quality will oscillate around some equilibrium level.

The possibility of restoration significantly affects the optimal manage-

¹This is simply the dynamic analog to the familiar static efficiency condition that marginal benefit equals marginal cost.

ment of a resource. Restoration typically introduces nonconvexities in cleanup costs, upsetting the simple interior solution just described. Most commonly, restoration efforts have high fixed costs, introducing economies of scale. For example, one method of cleaning up a hazardous waste site is to haul the soil away and incinerate it, in which case the costs vary little with the concentration of the contaminant in the soil. Similarly, there are high fixed costs involved in scraping zebra mussels from a water intake pipe or hauling hazardous waste to a treatment facility. Or the source of the nonconvexity may be institutional: establishing a regulatory regime, such as a ban on fishing in Georges Bank, may entail significant political costs.

Given a nonconvexity, optimization is a more complex process, and cleanup may proceed in a jerky fashion. The mundane example of desk mess illustrates the process. A desk gets messier and messier, until restored with a sudden burst of cleanup activity. A similar cycle characterizes cleanup efforts for many environmental problems. Rather than preserve a steady state by achieving zero net flow of new environmental bads, we periodically reduce the stock. Examples include capping a landfill, dredging a harbor, hauling hazardous waste to a permanent treatment center, or clearing mussels from the intake pipe of a Great Lakes power plant.

In this paper, we develop a general model of the optimal management of natural resources when restoration (with economies of scale) is an option. As a benchmark, we first consider the conventional case, where the stock cannot be cleaned up directly, so that abatement is the only feasible strategy. We then introduce the possibility of restoration. We start with the simple case in which the flow of pollution is exogenous, and only restoration is possible. We show that this case is formally equivalent to an optimal-inventory problem, and the optimal restoration policy has the familiar interpretation of an (S, s) policy.

The focus of our analysis is on the complete SFQ case, in which flows are controllable and restoration is feasible. We find that when both restoration and abatement are possible, the optimal policy employs both strategies, and that neither strategy takes the form it would in the absence of the other. The same lesson applies when it is physical or human capital whose quality is being promoted.

The next section reviews the relevant economics literature. Section 3 introduces the basic model, and formally defines our notions of abatement and restoration. Section 4 develops the theoretical results, analyzing the optimal policies in the three scenarios described above: the abatement-only case, the restoration-only case, and the joint abatement-restoration case. Section 5 illustrates these results with examples of real-world SFQ problems.

It discusses beach contamination, the storage and shipment of atomic wastes, zebra mussel control, municipal solid waste management, and the handling of hazardous waste at Harvard University. It also suggests how our model can be applied to the management of physical and human capital. Section 6 concludes.

2 Previous literature

The literature on the economics of pollution control has focused almost exclusively on the economics of emissions reduction. This emphasis on abating the flow rather than cleaning up the stock leaves restoration to the engineers and to the real world, which has seen many wastes removed, landfills capped, and toxic waste sites restored.²

Nonetheless, a review of the literature on stock pollutants provides useful context for this paper.³ Several models (Keeler, Spence, and Zeckhauser 1971; Plourde 1972; Plourde and Yeung 1989; Smith 1972) analyze the problem in a general-equilibrium setting, attending to the trade-off between pollution control and consumption or capital accumulation. These models generally assume strict convexity in abatement costs, leading to steady-state interior solutions in which the flow of new pollution just offsets the natural decay of accumulated pollution. Similar results obtain in models that frame the management of a particular pollutant as an optimal control problem, with abatement as the choice variable (Falk and Mendelsohn 1993). The steady-state solution is found by equating the marginal cost of abatement with the present-value marginal benefit of pollution reduction.

Caputo and Wilen (1995) address the problem of cleaning up a fixed pollution stock. Their example is hazardous waste. They consider restoration, as we define it, although they assume that cleanup costs are convex. They find that the optimal solution stops short of complete cleanup (and lets nat-

²Two factors may help to explain this treatment. First, the authors of these studies have typically examined cases in which restoration is essentially infeasible (e.g., the accumulation of greenhouse gases in the atmosphere). Second, previous models of stock pollution have employed optimal control theory to solve the complex dynamic optimization problem. In doing so, however, authors have typically viewed the flow of pollution (i.e., emissions) as the choice variable, and the stock as the state variable. In a sense, the methodology imposed the dichotomy between stocks and flows, and precluded the consideration of restoration efforts, which involve direct reductions in the stock.

³A pure stock (as opposed to flow) pollutant is one that persists indefinitely in the environment. More generally, we use the term “stock” or “accumulating” pollutant to refer to a pollutant for which new flows outpace decay. Examples include atmospheric greenhouse gases and hazardous waste.

ural degradation finish the process), as long as when pollution approaches zero so does its marginal damage.

Phillips and Zeckhauser (1998) argue that restoration efforts typically involve “destination-driven costs,” an extreme case of economies of scale. Destination-driven costs depend primarily on the ultimate level of environmental quality, rather than the initial level of environmental quality (or the amount of cleanup needed). In such cases, optimal cleanup does not follow the usual “marginal benefit equals marginal cost” prescription. Phillips and Zeckhauser, however, treat the restoration problem in a static, one-time-only setting. Their analysis addresses the question: “The pollution is there, or possibly accumulating; what shall we do?” This paper begins from the premise that pollutants will continue to arrive in the future.

3 Model framework

In this section, we construct a model of environmental degradation and amelioration that allows us to analyze optimal abatement and restoration policies. At this level of abstraction, we have in mind a generalized environmental resource with a “quality” level that changes over time. In the case of accumulating waste, for example, the quality might be measured by the volume of waste: the smaller the amount, the higher the level of environmental quality.⁴ We represent the quality of the environmental resource at time t by a real number x_t . Larger values of x_t represent more desirable states. We normalize the initial quality level to be equal to zero, so that $x_0 = 0$, and we shall be working mostly with negative values for x .⁵

Apart from any efforts of a resource manager, two processes acting in opposite directions affect the level of environmental quality: ongoing damage to the resource and natural recovery processes, such as the decay of the accumulated pollution stock. To capture both effects, we model the process of injury or damage to the resource as a random variable with drift. The amount of damage incurred up to time t , denoted by z_t , is assumed to follow a Brownian motion with drift rate $\mu > 0$, variance rate σ^2 , and $z_0 = 0$.

⁴Note that we use “quality” to denote the state of the resource: for example, it may measure how much pollution has accumulated. How the quality of a resource is *valued* will be captured in the utility function.

⁵Of course, the state x can be positive as well – the state evolves according to a stochastic process, and may improve as well as deteriorate. But since we are concerned primarily with deterioration, and we normalize $x = 0$ to be our starting point, negative values for x will be of primary interest.

Hence, damages evolve according to

$$(1) \quad z_t = \mu t - \sigma w_t,$$

where w_t follows a standard Brownian motion. Intuitively, μ can be thought of as the “average” rate of injury to the resource: for example, the average flow of pollution minus the natural decay of existing pollution. The random term in equation (1) captures random variations in the processes of damage and natural recovery. Accumulated damages reduce environmental quality x_t . Absent measures that reduce the flow of damage or restore the resource, quality at time t is $x_t = -\mu t + \sigma w_t$.

We assume that society’s benefit from the resource at any point in time depends only on the level of environmental quality. Thus, at time t society derives a flow of utility $u(x_t)$ from the availability of the resource.⁶ We assume that the social rate of time preference is $\alpha > 0$. We further assume that the utility function has the following properties.

Assumption 1 *The utility function u is twice continuously differentiable, with $u < 0$, $u' > 0$, $u'' < 0$, and u' unbounded above. Furthermore, $E_x [\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt]$ is finite for all x .*

Note that when environmental quality is less than zero, so is utility; the utility function can be thought of as the negative value of a convex loss function.⁷

3.1 Abatement Policies

We define abatement as a reduction in the rate of injury to the resource. In our SFQ framework, abatement corresponds to flow control. End-of-pipe controls on pollution emissions and changes in the production process that reduce pollution are both forms of abatement.

⁶We ignore issues such as population growth or changes in income, which could make the utility function time-dependent. With a growing population, for example, one might scale the utility function to the size of the population, so that the absolute value of the (negative) utility associated with a given level of environmental damage (negative quality) would increase over time. If abatement costs remained constant, the optimal level of abatement at a given level of environmental quality would increase over time. On the other hand, we might expect that abatement costs and the drift rate μ might be greater for a larger population producing more waste.

⁷Having a utility function with values less than zero simplifies notation, while letting us still discuss results in terms of “net benefits,” and so on. Of course, adding a suitable constant term to the utility function would make the values of the utility function positive over most of its range, and would not affect the results.

The resource manager’s problem is to determine the optimal abatement rate.⁸ Abating at rate $a(x_t)$ reduces the mean rate of injury from μ to $\mu - a(x_t)$. We assume that the abatement rate cannot exceed some finite ceiling \bar{a} . That is, infinite abatement is assumed to be impossible, whether for physical reasons (it cannot be achieved by existing technologies) or budgetary ones (the cost of infinite abatement exceeds the resources available). This assumption is invoked for technical reasons, and is not necessary for our results.⁹ In our discussion of optimal policies, in Section 4 of the paper, we also assume that the ceiling \bar{a} is greater than the mean flow rate μ . Thus the manager can, if she wishes, maintain a steady state in expectation by setting $a = \mu$.

An *abatement policy* is a mapping $a : \mathfrak{X} \mapsto [0, \bar{a}]$ from the set of real numbers (the possible values of the state x) to the interval $[0, \bar{a}]$ (the feasible levels of abatement). Thus an abatement policy specifies the abatement level as a function of the state x .¹⁰ Under an abatement policy a , the state of

⁸We use the term “manager” to refer to the person (or department) responsible for determining abatement and restoration. Such a “manager” may not actually control the resource directly, but rather may be a regulator who determines the ideal abatement and/or restoration policy and then employs a policy instrument (direct controls, standards, taxes, etc.) to get that policy achieved.

⁹The assumption of a ceiling on abatement simplifies the proof that an optimal abatement policy exists. Somewhat loosely, the technical challenge is that allowing abatement to be unbounded could, in principle, lead to situations where no policy is optimal (i.e., one can always do better by choosing a higher level of abatement). Requiring that abatement be bounded by a constant, as we have done here, ensures the existence of an optimal abatement policy.

In fact, for the problem considered here, an optimal abatement policy can be shown to exist even if we allow abatement to be unbounded. In the abate-only case, the concavity of the value function and the convexity of the abatement cost function are sufficient to ensure that infinite abatement will never be optimal. In the abate-and-restore case, in which the value function has a convex region, it is also possible (but more complicated) to show that the second derivative of the abatement cost function is greater than the second derivative of the value function, at least for sufficiently high rates of abatement. This rules out the possibility of infinite abatement being optimal in the abate-and-restore case as well.

In any case, the assumption of a ceiling on abatement has little practical import, because the ceiling can always be set high enough that the probability it binds is vanishingly small.

¹⁰We are implicitly confining the manager’s problem to that of choosing the optimal *stationary* policy, in which the choice of abatement is a function of the state. An alternative approach would be to make the manager’s problem one of choosing an optimal stochastic process $\{a_t\}$ measurable with respect to the filtration generated by $\{w_t\}$. However, one can show that such an optimal process can be produced by letting $a_t = a(x_t)$. Thus our assumption that the manager chooses the best among the set of stationary policies does not affect the practical implications of the analysis.

the resource evolves according to $dx_t = (a(x_t) - \mu)dt + \sigma dw_t$.

The resource manager faces the classic trade-off between the benefits of higher environmental quality and the costs of achieving it.

Assumption 2 *The abatement cost function $c : [0, \infty)$ is twice continuously differentiable with $c \geq 0$, $c(0) = 0$, and $c'' \geq \epsilon$ for some $\epsilon > 0$.*

Given an abatement policy a and an initial state x , the expected discounted utility over an infinite horizon can be written as

$$E_x^a \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right].$$

We assume that the manager wishes to choose an abatement policy that maximizes this expectation simultaneously for all x .

3.2 Restoration Policies

Next, consider a setting in which abatement is no longer an option, but instead it is possible to restore the resource. For example, the manager of a polluted stretch of coastline may be unable to prevent an ocean vessel from emptying its bilge offshore, but can replenish the sand on the beach.¹¹ Similarly, the causes of beach erosion or silt build-up in a harbor may be natural and outside the manager's control; restoration may be the only remedy available.

In our SFQ framework, restoration corresponds to an "all-at-once" cleanup that affects the stock of environmental quality directly, rather than by reducing the flow of pollutants. In particular, from any state x_t , the manager may choose to restore the resource to the state $x = 0$.¹² Letting τ_i denote the i th time at which the resource is restored, the state at time t is given by $x_t = -\mu t + \sigma w_t - \sum_{\{i|\tau_i < t\}} x_{\tau_i}$.

We assume that there is a positive fixed cost of restoration, and zero marginal costs. Thus the restoration cost is independent of the starting-point of restoration.

¹¹Thus exogeneity may be due to jurisdictional boundaries rather than intrinsic properties of the resource. The degree of institutional separation is clearly a function of scale: beach pollution that is exogenous from the point of view of a local coastal planner may be endogenous in the context of an international ocean-dumping protocol.

¹²A more general framework would explicitly model the manager's choice of how much to restore – that is, how far to clean up the resource. (For example, there might be more than one possible restoration method.) But in general, such a problem will yield a single optimal destination, so little is lost by our assumption of one restoration technology. Our choice of $x = 0$ as the destination, meanwhile, is merely a normalization and does not affect the results.

Assumption 3 *The cost of restoring quality from any state x_t to $x = 0$ is independent of x_t and is given by $C > 0$.*

While such “destination-driven” costs are an extreme form of cost non-convexity, the basic results of the model still hold under less extreme forms of scale economies, in which restoration involves constant or increasing marginal costs as well as a fixed cost.¹³

A restoration policy is characterized by a measurable closed subset R of \mathfrak{R} . Under a restoration policy R , restoration occurs whenever the state x_t occupies the set R . That is, the resource is restored whenever its quality falls to a certain point.¹⁴

Given a restoration policy R and an initial state x , the infinite horizon expected discounted utility can be written as

$$E_x^R \left[\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt - \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C \right].$$

We assume that the manager wishes to choose a restoration policy that maximizes this expectation simultaneously for all x .

3.3 Combined Abatement–Restoration Policies

Finally, let us consider a setting in which the manager can affect both stocks and flows independently. That is, abatement and restoration are both feasible. The state of the resource evolves according to $x_t = \int_{s=0}^t (a(x_s) - \mu) ds + \sigma w_t - \sum_{\{i|\tau_i < t\}} x_{\tau_i}$. Given a combined abatement–restoration policy (a, R) and an initial state x , the infinite horizon expected discounted utility can be written as

$$E_x^{a,R} \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt - \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C \right].$$

The manager’s objective is to choose a combined abatement–restoration policy that maximizes this expectation simultaneously for all x .

¹³Suppose that the restoration cost is given by $C(x) = F + \gamma(x)$, with $\gamma(x)$ a convex variable-cost function. Unless $J(0) - J(x) < C(x)$ for all x , restoration will be optimal for at least one state x . If we now let \underline{x} denote the largest “trigger” (that is, the highest value of x at which restoration is optimal), then the system will evolve much as in the case with only a fixed cost. A difference will be that the abatement rate need not go to zero as the state approaches \underline{x} .

¹⁴Again, we are implicitly confining the manager’s choice set to the set of stationary policies. But as in the abatement case, a stationary policy would still be optimal if we allowed for a broader class of policies. See note 10.

4 Optimal policies

In this section we study optimal policies in the three settings described above. In the first case, only abatement is possible. This provides a benchmark for the results when restoration is feasible. In the second case, flows are exogenous, and restoration is the only option. This case allows us to focus on the dynamics of optimal restorations.¹⁵ We devote the bulk of our attention to the third case, in which both abatement and restoration are feasible. Proofs of all theorems may be found in Appendix A.

4.1 Abatement Policies

Let J be the optimal value function:

$$(2) \quad J(x) = \sup_a E_x^a \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right],$$

where the supremum is taken over abatement policies. $J(x)$ represents the maximal present value of the future stream of net benefits (utility minus cost) under the optimal policy, starting from state x .

Theorem 1 *Let Assumptions 1 and 2 hold. Then,*

1. $J < 0$ and $J(x)$ is finite for every x ;
2. J is twice continuously differentiable;
3. J satisfies

$$(3) \quad \sup_{q \in [0, \bar{a}]} \left(\frac{\sigma^2}{2} J''(x) + (q - \mu) J'(x) - \alpha J(x) + u(x) - c(q) \right) = 0,$$

for all x ;

4. $J' > 0$;
5. $J'' < 0$;
6. J' is unbounded above;

¹⁵The reader will note that in formal terms the restoration-only case is simply a special case of the abatement-and-restoration case, with the ceiling $\bar{a} = 0$. We take advantage of this fact to simplify our proofs. For expositional clarity we discuss the restoration-only case separately, before considering both abatement and restoration together.

7. there is a function $a^* : \mathfrak{R} \mapsto [0, \bar{a}]$ such that for every x , $a^*(x)$ uniquely attains the supremum in equation (3); and
8. a^* is an optimal policy.

From Theorem 1, the optimal value function J satisfies the Hamilton-Jacobi-Bellman equation (equation (3)).¹⁶ Thus the value $a^*(x)$, which attains the supremum, constitutes an optimal abatement decision at state x . The following theorem describes this optimal abatement policy.

Theorem 2 *Let assumptions 1 and 2 hold. Then,*

1. there exists a state \hat{x} such that a^* is decreasing on (\hat{x}, ∞) and $a^*(x) = \bar{a}$ for all $x \leq \hat{x}$;
2. $\lim_{x \rightarrow \infty} a^*(x) = 0$; and
3. there exists a state x^* such that $\mu < a^*(x)$ for $x < x^*$ and $\mu > a^*(x)$ for $x > x^*$.

Theorem 2 establishes that as environmental quality falls, the optimal rate of abatement increases.¹⁷

The optimal abatement policy $a^*(x)$ has a familiar economic interpretation.¹⁸ Most of the terms in equation (3) are functions only of the state x and not of the abatement decision q ; since the state is not a choice variable, a value q attains the supremum if and only if it also attains the supremum $\sup_{q \in [0, \bar{a}]} f_x(q)$ of a function $f_x(q) = qJ'(x) - c(q)$. Intuitively, $f_x(q)$ can be understood as the expected net benefit from abatement q . The second term, $c(q)$, represents the cost of abatement q . The first term, $qJ'(x)$, represents the rate at which the value function increases. This corresponds roughly to the expected *benefit* from abating at rate q , taking into account present

¹⁶Equation (3) is a close relative of the Bellman equation commonly used in recursive treatments of economic problems. For a rigorous development of recursive methods and their application to a range of problems in economics, including the (S, s) inventory problem we discuss below, see Stokey and Lucas (1989).

¹⁷Of course, at some point the abatement policy will be constrained by the “ceiling” \bar{a} .

¹⁸This and subsequent heuristic explanations of the results in economic terms are less technically rigorous than the theorems themselves. These discussions aim to connect our formal results to more familiar results from economics. Intuition’s gain is occasionally rigor’s loss.

and future utility.¹⁹ Theorem 2 therefore states that the optimal abatement policy at any given time picks the abatement rate to maximize the resulting “expected net benefit,” assuming it is feasible. If not, the optimal policy is to abate as much as possible – that is, \bar{a} .

Given the concave utility function, the marginal loss of utility from further damages increases as environmental quality deteriorates. Thus the marginal benefit from abatement increases as x diminishes. The marginal cost of abatement, however, does not depend on environmental quality, but only on the amount of the flow abated. With marginal benefit increasing as x falls, and marginal cost staying constant, the optimal rate of abatement rises accordingly, as Theorem 2 shows.

The abatement function $a^*(x)$ described in Theorem 2 can be compared to the optimal policy in a deterministic world, in which the unabated flow of damages never deviated from its mean rate μ . In such a world, the optimal abatement policy would maintain a level of environmental quality, x^* , where abatement exactly equalled the rate of damage: $a^*(x^*) = \mu$. For $x > x^*$, $a^*(x) < \mu$, meaning that environmental deterioration would outpace abatement efforts, and the state of the environment would decline towards x^* . For $x < x^*$, on the other hand, $a^*(x) > \mu$ and abatement would more than compensate for damages, yielding a continuous improvement in environmental quality. Starting from any level of environmental quality, the optimal policy would reach an equilibrium at a quality level x^* satisfying $J'(x^*) = c'(a^*(x^*)) = c'(\mu)$, where the marginal cost of fully abating new pollution (achieving a zero net flow) just equalled the marginal net benefits from doing so. This result, of course, agrees with standard results from the optimal-control literature discussed in Section 2.

Introducing uncertainty changes the realized outcomes, but the optimal policy can still be understood intuitively in terms of the “target” quality level x^* , where the rate of abatement that maximizes *expected* net benefits exactly offsets the *expected* flow of damages. We call x^* an *expectation equilibrium*, meaning that it is an equilibrium in expectation: at x^* , abatement equals the expected flow, and thus in expectation the quality level will remain there. Moreover, x^* is stable. Starting from x^* , a “high” flow of damages (i.e., at a rate greater than μ) will depress environmental quality below x^* .

¹⁹Heuristically, for a given marginal change in the state dx , the resulting change in the value function would be $J'(x)dx$. We can think of abatement q , carried out over an infinitesimal time period of duration dt , as yielding a marginal improvement in the state due to abatement $dx = qdt$. We can think of $(qdt)J'(x)$ as the change in the value function (over an infinitesimal period of time). Dividing through by dt yields the rate of change in the value function, $qJ'(x)$.

In response, abatement will increase, so that $a^*(x) > \mu$. Thus the quality level will return to the target x^* in expectation. A “low” flow of damages will raise quality above its target level, leading to a slackening of abatement efforts and a tendency back toward x^* .

4.2 Restoration Policies

In many cases of interest in the real world, abatement is not possible, but the resource manager can periodically restore environmental quality. The prime result of this section, Theorem 3, shows that in such cases the optimal policy is to allow environmental quality to decline until the state falls to a certain level, at which point the manager undertakes restoration. Environmental degradation then resumes, environmental quality starts to decline (albeit stochastically), and the cycle repeats.²⁰

Let J be the optimal value function:

$$J(x) = \sup_R E_x^R \left[\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt - \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C \right],$$

where the supremum is taken over restoration policies.

Theorem 3 *Let Assumptions 1 and 3 hold. Then, there exist states \underline{x} and x^\dagger with $\underline{x} < x^\dagger$ such that the following statements hold:*

1. $J < 0$ and $J(x)$ is finite for every x ;

²⁰In some real-world cases, restoration is done only once, after which the process of environmental degradation stops. While our model does not explicitly consider such one-time-only restorations, the implications of the model are reasonably clear for such cases: we abate at the margin (if possible) and then restore all at once. In other words, we maximize up until the first restoration, and then eliminate new flows after that. As discussed above, the costs of eliminating further flows may be treated simply as part of the restoration cost. In order for one-time restoration to be optimal, however, there would presumably need to be a nonconvexity in benefits, or at least very high marginal damages from the first few units of degradation.

Note that some cases which might appear to involve “one-time only” restorations may, in fact, simply be combined abatement-restoration policies in which the abatement policy is designed to maintain a fairly high level of environmental quality, implying that the probability that the resource quality declines far enough to trigger restoration is very small. For example, when we clean up a river, we may not expect to restore it again; but this expectation may be implicitly based on our attempts to maintain a high river quality, rather than on an absolute cessation of the flow of damages. In the case of a river restoration, one could argue that the possibility remains – however slight – that unforeseen changes in water currents or waste generation will require future restorations. See the discussions of combined abatement-restoration policies in Section 4.3 below.

2. $J(x) = J(0) - C$ for all $x \leq \underline{x}$;
3. $R = (-\infty, \underline{x}]$ defines an optimal restoration policy;
4. J is twice continuously differentiable on (\underline{x}, ∞) ;
5. J is continuously differentiable;
6. J satisfies

$$\frac{\sigma^2}{2}J''(x) - \mu J'(x) - \alpha J(x) + u(x) = 0,$$

for all $x > \underline{x}$;

7. $J' \geq 0$;
8. $J'(x) > 0$ for $x \in (\underline{x}, \infty)$;
9. $J''(x^\dagger) = 0$;
10. $J''(x) > 0$ for $x \in (\underline{x}, x^\dagger)$; and
11. $J''(x) < 0$ for $x \in (x^\dagger, \infty)$.

The optimal restoration policy is thus defined by a “trigger level” \underline{x} . When the state reaches this lower bound, the resource is restored up to $x = 0$. The periodic restoration of a resource at a fixed cost is thus a form of the “optimal stopping problem.”²¹ Indeed, the optimal restoration of an environmental resource resembles the familiar solution to the classic inventory problem: a profit-maximizing firm will follow an (S, s) rule in managing its inventory, drawing its stock of goods down until some level s is reached and then replenishing the inventory up to the level S (Arrow, Harris, and Marschak 1951; Scarf 1960). The “inventory” in the restoration case is the level (stock) of environmental quality, and a restoration corresponds to a replenishment of inventory. Despite this resemblance, however, our control-of-quality problem differs from the inventory problem in a key respect. The cost in both cases comes from restoring the stock, but the source of the payoff differs. A firm managing its inventory receives an immediate payoff

²¹To use Rust’s (1987) term, the restoration problem is an example of a *regenerative optimal stopping* problem. Here, a restoration “regenerates” the state of the resource to $x = 0$. Rust contemplated the optimal replacement of bus engines; we discuss the applications of our model to machine replacement and other aspects of investment in physical capital in Section 5, below.

from a change in stock. In our control-of-quality problem, the immediate payoff derives from the size of the stock.²²

How far the quality of the resource must fall before restoration becomes optimal depends on the cost of restoration C . The following Theorem establishes that as C increases, the trigger level \underline{x} decreases. Intuitively, a higher restoration cost increases the incentive to delay restoration.

Theorem 4 *Let $\underline{x}(C)$ denote the trigger level for a given restoration cost $C > 0$. Then $\underline{x}(C)$ is decreasing in C .*

4.3 Combined Abatement–Restoration Policies

We now consider the complete SFQ problem, where the manager can use both abatement and restoration. The results of this section show that when both clean-up methods are available, both are employed. Moreover, the availability of restoration affects the optimal abatement policy, and the possibility of abatement alters restoration. Theorems 5 and 6 describe the form of the optimal policy and derive this “non-separability” result. We then explore how the optimal policy changes with the mean flow of new damages.

Let J be the optimal value function:

$$J(x) = \sup_{a,R} E_x^{a,R} \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt - \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C \right],$$

where the supremum is taken over pairs of abatement and restoration policies.

Theorem 5 *Let Assumptions 1, 2, and 3 hold. Then, there exist states \underline{x} and x^\dagger with $\underline{x} < x^\dagger$ such that the following statements hold:*

1. $J < 0$ and $J(x)$ is finite for every x ;
2. $J(x) = J(0) - C$ for all $x \leq \underline{x}$;
3. J is twice continuously differentiable on (\underline{x}, ∞) ;
4. J is continuously differentiable;

²²Balcer (1973) adds a flow-control instrument to a traditional optimal inventory problem to analyze comparative statics. Specifically, he makes advertising a choice variable that positively affects consumer demand. Hence, the level of advertising must be considered jointly with the selection of the (S, s) bounds.

5. J satisfies

$$(4) \quad \sup_{q \in [0, \bar{a}]} \left(\frac{\sigma^2}{2} J''(x) + (q - \mu) J'(x) - \alpha J(x) + u(x) - c(q) \right) = 0,$$

for all $x > \underline{x}$;

6. $J' \geq 0$;

7. $J'(x) > 0$ for $x \in (\underline{x}, \infty)$;

8. $J''(x^\dagger) = 0$;

9. $J''(x) > 0$ for $x \in (\underline{x}, x^\dagger)$;

10. $J''(x) < 0$ for $x \in (x^\dagger, \infty)$;

11. there is a function $a^* : (\underline{x}, \infty) \mapsto [0, \bar{a}]$ such that for every $x \in (\underline{x}, \infty)$, $a^*(x)$ uniquely attains the supremum in equation (4); and

12. letting $R^* = (-\infty, \underline{x}]$, the pair (a^*, R^*) is an optimal policy.

This theorem identifies two states: the “trigger” level \underline{x} and an inflection point x^\dagger . The manager restores the resource when environmental quality reaches or falls below \underline{x} , as in the case when only restoration is possible. The inflection point x^\dagger marks a point of transition in the optimal value function J : below x^\dagger , J is convex; above it, J is concave.²³

The next theorem describes the optimal combined abatement-restoration policy (a^*, R^*) defined by Theorem 5, and compares the optimal abatement rate with and without the possibility of restoration. To aid the comparison, let \tilde{J} and \tilde{a} denote the optimal value function and abatement policy, respectively, when restoration is not an option. (Thus \tilde{J} and \tilde{a} are defined as in Theorems 1 and 2.)

²³Theorem 6 shows that when restoration is feasible, optimal abatement is not a monotonic function of quality, and reaches a peak at the inflection point x^\dagger . The inflection point x^\dagger may therefore be worth explaining in more detail. Recall that when only abatement is possible, the value function is concave, reflecting the concavity of the underlying utility function. Under a restoration regime, however, the value function is constant below \underline{x} , since the restoration always returns the state to $x = 0$ at a fixed cost. Because J is differentiable, its slope at \underline{x} is zero. Above \underline{x} , J is increasing. In some region just above \underline{x} , therefore, $J(x)$ must be convex. This region is bounded by \underline{x} below and x^\dagger above.

Theorem 6 *Let Assumptions 1, 2, and 3 hold. Then,*

1. $J' < \tilde{J}'$;
2. a^* is increasing on $\{x \in (\underline{x}, x^\dagger) | a(x) \neq \bar{a}\}$ and decreasing on $\{x \in (x^\dagger, \infty) | a(x) \neq \bar{a}\}$; and
3. for each state $x \in (\underline{x}, \infty)$, either $a(x) < \tilde{a}(x)$ or $a(x) = \tilde{a}(x) = \bar{a}$.

The first assertion states that the derivative of the value function – the marginal increase in the present value of net benefits as the resource’s state improves – is everywhere less in the abate-and-restore case than in the abate-only case. The feasibility of restoration raises the value function everywhere, since its absence represents a constraint on the resource manager. But the value function increases more at low levels of quality, where restoration is imminent, than at high levels of quality, where restoration is more distant.

Assertions 2 and 3 establish a *non-separability* result: the possibility of restoration alters the optimal path of abatement. Recall from Theorem 2 that when only abatement is possible, the optimal level of abatement always increases as environmental quality worsens. Assertion 2 shows that this monotonic relationship fails to hold when restoration is an option. As environmental quality declines from its high initial level, the optimal rate of abatement first increases, but then slows as the state approaches the “trigger” level \underline{x} from above. This result is intuitive. Restoration reduces the incentive to abatement the flow of damages: because the quality of the resource will be restored to zero when the point \underline{x} is reached, abatement will yield benefits only until the moment of restoration, as opposed to forever in the abate-only case. As the trigger point is approached, restoration becomes ever more imminent; hence the benefits from abatement diminish.

By Assertion 3, the optimal level of abatement when restoration is possible is everywhere weakly lower than in the abate-only case. This result follows closely from Assertion 1. When restoration is possible, the present value of net benefits (i.e., the value function $J(x)$) increases more slowly as the state improves. Loosely speaking, the marginal gains from abatement are lower. Hence less abatement is performed.

These results are illustrated by Figure 1, which shows the optimal abatement policies with and without the possibility of restoration, for a particular set of functional forms and parameter values.²⁴ Note that when restoration is

²⁴The functional forms and parameter values used for all figures are described in Appendix B.




Figure 1: Optimal abatement policies with and without the possibility of restoration.

feasible, the abatement rate falls all the way to zero at the restoration trigger point \underline{x} . This follows from the continuity of the first derivative of the value function, $J'(x)$, or the “smooth-pasting” condition (Krylov, 1980). Because restoration will take place at any $x < \underline{x}$, the value function $J(x)$ is constant below \underline{x} , and thus its derivative is zero. The smooth-pasting condition implies that $J'(x)$ goes continuously to zero as x goes to \underline{x} . Since abatement is chosen to maximize the “net benefits” function, $f_x(q) = qJ'(x) - c(q)$, abatement must go to zero as x goes to \underline{x} .

The smooth-pasting condition is intuitive. When the resource quality is at or below the restoration point \underline{x} , abatement brings zero benefits, because the state will be restored to $x = 0$ anyway.²⁵ The smooth-pasting condition implies that as the quality of the resource approaches the trigger point \underline{x} , the benefits from abatement must decline smoothly to zero. If the benefits decline to zero, however, the costs must do the same. Thus abatement must go to zero as well.

Figure 2 shows the corresponding value functions J and \tilde{J} , along with the value function in the restore-only case. Because the absence of restora-

²⁵Recall that $qJ'(x)$ can be thought of as the benefit from abating at rate q when the state is x .




Figure 2: Value functions for the abate-only, restore-only, and abate-and-restore cases.

tion or abatement as an option represents a constraint on the manager, the value functions in the abate-only and restore-only cases lie below the value function in the combined abate-and-restore case.

Figure 3 illustrates the optimal abatement functions for three values of the mean flow rate μ . The vertical axis measures the fraction of mean flow abated, and the horizontal axis measures resource quality x . If flows are low on average (the drift rate μ is small), the optimal abatement rate will rise above mean flow for some range of x , just as it does when abatement is the only option. In this case, $a(x) = \mu$ at two points: x^* and x^{**} .²⁶ The higher value, x^* , is a stable expectation equilibrium. Once the quality of the resource hits x^* , it will tend to oscillate around that level for small perturbations in the flow of damages. Small stochastic increases in the flow above

²⁶In fact, for some flow rate the optimal abatement function will be tangent to the horizontal line at μ . In this case, $a(x) = \mu$ at a single value of x , and then falls back to zero. The resulting “expectation equilibrium” x^* is then stable from the right but not the left. For $x > x^*$, abatement will be lower than μ and the state will tend to return to x^* . For $x < x^*$ abatement will also be less than μ , implying that in expectation the quality level will decline to the trigger level for restoration.

The flow rate leading to this “borderline case” is marked by the vertical dotted line in Figure 5 below.




Figure 3: Fraction of flow abated as a function of quality, for three flow rates.

μ will lead to countervailing increases in abatement, just as small decreases in the flow will induce abatement less than μ . Nonetheless, the possibility of restoration still affects the outcome. Since the optimal abatement rate is lower than if restoration were not an option, the expectation equilibrium occurs at a lower level of quality. Moreover, a large flow of damages can upset this equilibrium. In particular, if the resource quality falls below x^{**} , optimal abatement falls below the mean flow μ . If this happens, the quality level (in expectation) deteriorates steadily to the trigger level \underline{x} , whence the resource is restored.

The availability of restoration can also change the optimal abatement policy in a more fundamental way. Since $a^*(x)$ is no longer monotonic when restoration is possible, abatement need not rise above μ ; thus an expectation equilibrium may not exist. In this case, at *all* values of x above \underline{x} , abatement merely slows – but does not halt – the net flow of damages. Rather than maintaining quality at a certain level, the optimal policy lets damages accumulate over time until the trigger level is reached, and then restores the resource. This will occur if average flows are high (the drift rate μ is large). When the mean flow rate is high, the cost of offsetting it with abatement is high as well. At the same time, high flows mean that restoration will be

more frequent, on average, so that damages will persist in the environment for a shorter period of time before being cleaned up. Hence at higher flow rates, restoration is employed more relative to abatement.

Thus, when flows are low, abatement exceeds the average flow over some values of x , producing an expectation equilibrium at the point where abatement equals the mean flow. Although restoration will occur if quality falls far enough, we are quite unlikely ever to reach that region. At higher flow rates, abatement is always less than the average flow, and no expectation equilibrium exists.²⁷ For low flow rates, then, we may say that abatement is the “principal strategy.” For high flow rates, restoration is the principal strategy. But in both cases, the presence of one strategy affects the use of the other.

Figure 4 shows more directly how the importance of abatement relative to restoration varies with the flow rate. The horizontal axis measures the flow rate. The vertical axis measures the time-averaged rate of abatement as a fraction of the flow of damages, or (equivalently) the fraction of total damages that is cleaned up by abatement rather than restoration. Thus for a given flow rate, the height of the curve represents the fraction of cleanup due to abatement. The remainder of the cleanup, from the curve to the top of the graph, is due to restoration. (For example, at a flow rate of 3, approximately 20% of the total flow of damages is cleaned up by abatement, with the remaining 80% cleaned up through periodic restorations.) If flows are low, abatement offsets flows completely (in expectation) almost everywhere, so that restoration occurs with very small probability and virtually all of the damages are cleaned up through abatement. As flows increase, the fraction of flow abated drops dramatically.

Figure 4 also illustrates how the existence of an expectation equilibrium depends on the flow rate. For flow rates to the left of the dotted vertical line, an expectation equilibrium exists. At higher flows, to the right of the line, no expectation equilibrium exists.

Whether cleanup relies more on restoration or abatement determines how the quality of the resource varies over time. Figure 5 plots the frequency distribution of states for the same three flow rates as in Figure 3. When flows are low, an expectation equilibrium is achieved. States close to this

²⁷The focus of our discussion is on the fraction of flow abated, which clearly reaches a higher peak at low flow rates than at high ones. As Figure 3 shows, the peak of the abatement function may also be higher in *absolute* terms when flows are low. At a flow rate of 3, for example, abatement rises to roughly one quarter of the flow rate, or 0.75 in absolute terms. At $\mu = 0.9$, on the other hand, abatement reaches over twice the flow rate, or 1.8 in absolute terms.




Figure 4: Fraction of cleanup due to abatement, as a function of the mean flow rate μ .

equilibrium level are much more common than other states, as shown by the peak of the frequency distribution. At somewhat higher flow rates, no expectation equilibrium exists, and restorations occur more frequently. As a result, high-quality states become more common. Nonetheless, since less abatement occurs when the quality of the resource is high, the most common states are those somewhat above the restoration point, producing a peak in the frequency distribution. At a high flow rate, restoration becomes more important relative to abatement. As a result, all states between the initial quality level and the restoration point occur with roughly equal frequency.

Figure 5 also demonstrates a seemingly paradoxical result: in the full SFQ case, when both restoration and abatement are possible, the average quality of the resource turns out to be higher with a high flow of damages than with a low (or medium) flow.²⁸ The reason is that restorations are more frequent when the flow of damages is high, so that high-quality states are more common. Although this result is interesting, it is not general. In Figure 5, abatement costs are high enough (relative to the cost of restoration) that

²⁸Of course, looking only at average quality ignores the costs of abatement and restoration. While the average quality of the resource may be higher with a high flow than with a low flow, the value function will always be *lower* with a high flow.




Figure 5: Frequency distributions of resource qualities (states) under optimal policies for three flow rates.

the expectation equilibrium in the low-flow case is at a fairly low level of quality (relative to the restoration trigger point). If marginal abatement costs were less, the expectation equilibrium would occur at a higher level of quality. If restoration were more expensive, the restoration trigger point would be lower. Thus the optimal policies, and the quality of the resource over time, depend on the relative costs of abatement and restoration as well as the flow rate.

5 Applications to the environment, and to physical and human capital

In this section, we explore the results of our theoretical models in a range of real-world cases. We start by considering examples of SFQ problems in natural resources and the environment. We then show how our model applies to the optimal management of two more conventional resources, physical and human capital.

5.1 SFQ problems in the environmental arena

We now turn to real-world examples illustrating restoration with and without the possibility of abatement.

5.1.1 Exogenous flows

When flows are exogenous, the resource manager only controls the stock of the resource. While this is a limiting case of the SFQ model, the restoration-only case has widespread application to the environmental arena. When it applies, the optimal policy contradicts the usual dictum to equate marginal benefits and marginal costs. For example, consider a stretch of ocean shoreline subject to seaborne pollution (medical waste, perhaps, or the tarry bilgewater from ocean-going tankers). The “conventional” model of stock pollution would suggest that the manager ought to strive for a steady-state level of quality, which implies daily cleansing. If there are economies of scale in clean-up, however, the quality of the shoreline should be allowed to vary over time: pollution accumulates, the beach is cleaned up all at once, pollution builds up again, and the cycle repeats. The higher the restoration cost, the greater the accumulation of waste before cleanup.

Atomic waste from nuclear power plants, as well as from atomic weapons production and disposal, medical centers, and so on, presents a similar pattern.²⁹ Hauling waste from a nuclear power plant involves expensive safety precautions, no matter how small the shipment. This makes frequent waste shipment uneconomic. Hence under an optimal policy, wastes accumulate, imposing increasing risks as they do, and then are periodically removed. The trigger level for shipment will increase with the cost of shipping relative to the period costs of damages from the accumulated waste.

5.1.2 Restoration and controllable flows

When both restoration and flow control are possible, we enter a full SFQ world. We showed above that the availability of restoration tempers abatement efforts. Moreover, if a steady-state target for environmental quality exists, it will be lower than if restoration were not feasible. We consider three examples: the removal of zebra mussels, the management of municipal solid waste, and the collection and shipment of hazardous waste generated in chemistry labs at Harvard University.

²⁹Of course, in actuality nuclear power plants in the United States have been storing waste on site as a “temporary” measure while awaiting a permanent disposal site (whose construction has been held up by geological uncertainty and political recalcitrance).

Zebra mussels. Zebra mussels (*Dreissena polymorpha*) are small freshwater mollusks, which clog water intake and distribution systems by adhering in large clusters to pumphouses, plumbing systems, and other pieces of equipment. The control of zebra mussels by power plants, water works, and other large users of water in the Great Lakes area of the United States is estimated to have cost from \$750 million to \$1 billion over the past decade.³⁰

The feasibility of preventing mussels from settling varies in different parts of a power plant, and the control strategies vary accordingly. The flow of zebra mussels is effectively exogenous in the pumphouses of power plants, where water diverted from lakes or rivers is stored and pumped into cooling water systems. Mussels grow on walls, debris screens, valves, and pumps, obstructing the flow of water. They can be scraped off once they settle, but they cannot be prevented from settling.³¹ Mechanical measures – physical scraping or “hydrolasing” with high-powered water hoses – involve high fixed costs: sending down a team of divers or even dewatering the pumphouse (thus shutting down the plant). An (*S,s*) policy is followed. Mussels are allowed to settle and grow, and periodically are removed. Removal is done every year or two in western Lake Erie, their densest habitat.

Inside the plumbing systems of power plants and water works, mussels are inaccessible to mechanical removal, but chemical removal is feasible. In such locations, both flow and stock controls are employed. For either, a chemical, such as chlorine, is injected into the water supply. Continuous low-level chlorination is an abatement policy. It kills juvenile mussels, and prevents their settlement. Periodic (annual or semi-annual) injections of high concentrations are a restoration (stock-control) strategy used to kill off adult mussels that have settled.³²

Municipal solid waste management. Next, consider the optimal management of municipal solid waste. The typical approach to managing garbage

³⁰Personal communication, Charles O’Neill, Project Director, National Zebra Mussel Information Clearinghouse, New York Sea Grant, March 24, 2000.

³¹Since these areas are open to lakes or rivers, chlorine and other chemical means would be ineffective in preventing settlement (and are ruled out anyway by environmental laws regulating water quality). Other measures, such as paints or surface materials designed to impede zebra mussel attachment, have not worked.

³²Continuous chlorination is typically effective enough that additional periodic treatments are unnecessary. This may be seen as an instance where flow control measures maintain a high-quality expectation equilibrium, and restorations are extremely rare. Nonetheless, if zebra mussel settlement were to occur (due, perhaps, to a breakdown in the chlorination regime), a one-time injection of chlorine at higher concentrations would be employed as a restoration measure.

is to build a landfill and slowly fill it over time, diverting some waste from it through recycling or composting. In our framework, diversion and source reduction reduce the flow of waste, and thus represent “abatement.” Capping a full landfill and building a park or a housing development on the site constitutes “restoration.”³³ In a typical scenario, waste diversion remains roughly constant over time, or changes only with changing preferences (i.e., a desire to increase levels of recycling) or price signals (e.g., land becomes more expensive, or recycled materials become more valuable).

Setting aside administrative obstacles or costs of policy adjustment, optimal waste management would have the abatement rate – the rate of waste diversion – vary over time. Consider the simplest case where the daily cost of a unit of waste is constant. When a landfill is first opened, diversion should be relatively high. That is because waste dumped early will be around for the landfill’s entire life. This implies that the discounted expected damages it imposes will be high relative to the damages from garbage arriving later. As the landfill nears capacity, waste diversion should slow, since the waste will impose damages only until the time of restoration, when the landfill is turned into a park.³⁴

Hazardous waste management at Harvard University. The management of hazardous waste in the chemistry department at Harvard University provides another illustration of an SFQ problem.³⁵ The department’s laboratories generate two types of waste, which are managed separately: a continual stream of solvents (used routinely in laboratory experiments), and more sporadic flows of a variety of toxic and reactive substances. Both types of wastes accumulate at individual laboratories. Hauling the wastes away for disposal – “restoration,” in this context – involves economies of scale.

³³With solid waste management, successive waves of accumulation and restoration take place on a series of dump sites, as opposed to the cyclical cleansing and soiling of a single resource. Our model can be extended to accommodate the multiple-site case by having the exposure costs and restoration costs rise as we move to successively more expensive landfills. Essentially, this would append results from the theory of nonrenewable resources to our models (Dasgupta and Heal, 1979; Hotelling, 1931). Abatement today would be influenced by the shadow price of future restorations.

³⁴This pattern – net dumping is speeded up as restoration approaches – plays out at the much smaller scale of the individual household. If garbage is picked up every Wednesday, then Tuesday not Thursday is the day to empty the refrigerator of leftovers.

³⁵The authors thank Henry Littleboy, Health and Safety Officer (in Harvard’s Faculty of Arts and Sciences Office of Environmental Health Services) overseeing hazardous waste management in the Chemistry Department, and Dr. Alan Long, Director of Laboratories, for their generosity in answering questions and providing information about hazardous waste management in the Harvard chemistry department.

How often should wastes be collected? Should they be shipped directly from individual labs, or consolidated in central collection depots first? How should waste generation be controlled?³⁶

The department's laboratories generate 200-250 gallons of solvents each week. Researchers collect the chemicals in 5-gallon cans, which are stored in "mini-main accumulation areas" (locked closets with safety equipment) near each lab. Once a week, an outside contractor makes the rounds in the department, collects 5-gallon cans from the mini-main outside each lab, and hauls them away. Harvard pays a fixed charge of \$23 or \$27 per 5-gallon container, depending on the type of solvent.

Solvents used to be picked up less frequently: once every two weeks in the mid-1990s, and once a month in the early part of the decade. Rather than being collected from individual labs, 5-gallon cans were collected in a storage area first. The vendor would then "bulk" the solvents, combining the contents of the cans into 55-gallon drums. Because such an approach requires less frequent visits by the contractor, it might be expected to cost less, per gallon of waste, than weekly collections. Records of actual costs, however, show that average disposal costs have gone *down* under the weekly collection regime. Very rough cost estimates suggest that monthly pick-ups cost approximately \$7 per gallon in 1990; bimonthly collections cost a little under \$6 per gallon in 1996; and weekly collections currently cost about \$5 per gallon. While this fall in cost appears to reflect increasing competition among vendors and better negotiating by Harvard, rather than diseconomies of scale, it suggests that cost savings from less frequent shipments would be slight.³⁷

Weekly collection is also preferable on other grounds. Solvents are highly flammable, and the potential for calamity from a storage room full of 5-gallon cans is much higher than it is from individual containers scattered around the department's buildings. Thus, as wastes accumulate, period utility falls sharply, while the per-unit costs of restoration are lessened only slightly (if at all) by waiting for larger shipments.

All other chemical wastes are hauled away in "lab packs": waste con-

³⁶Of course, chemical waste storage and disposal are heavily regulated by the Environmental Protection Agency. Existing regulations require the separate handling of different types of wastes (e.g., corrosives as opposed to pyroforics), and prohibit the storage of waste longer than ninety days. Nonetheless, the law provides some leeway for environmental managers to decide how best to dispose of their waste. At Harvard, for example, the prohibition on long-term storage is not a binding constraint, because limited storage space makes more frequent collection necessary.

³⁷This process continues: disposal charges for weekly collection of 5-gallon cans fell from \$30-35 to \$23-27 when a contract was signed with a new vendor in February, 2000.

tainers are collected from the labs, grouped by waste type, and then packed with similar wastes in larger drums. Rather than pouring the contents of the containers into the drums, technicians place the containers themselves into the drums and surround them with packing material. Thus, several 5-gallon cans of toxic flammable liquids might be collected in a 55 gallon drum. Flows of wastes shipped in lab packs are more variable than those of solvents; five 30-gallon drums of organic flammable liquids may be shipped out one week, and then none for the next month.

Lab pack hauls, like solvent shipments, now take place weekly. Until the end of 1999, chemical wastes were collected from individual labs and stored together in a “main accumulation area.” When that storage room was filled – every three to five weeks, on average – the waste would be shipped out. The building housing the main accumulation area, however, is slated for demolition. The storage room was cleared out this winter, and no replacement has been found. Because the waste is now collected directly from the individual labs, which lack storage space, lab packs are taken away each week.

While this arrangement may reflect the implicit value of laboratory space at Harvard, it will almost certainly raise the costs of waste disposal. Lab pack costs demonstrate considerable economies of scale, reflecting the underlying fixed costs of labor and transportation. A 55-gallon drum of corrosive flammable liquids costs \$320 to ship; a 30-gallon drum costs \$215, and a single 5-gallon container \$95. Similarly, shipping a 55-gallon drum of contaminated silica gel costs \$220, compared with \$45 for a 5-gallon can. In the recent weekly collections, many types of waste have been shipped in relatively small containers: 30-, 16-, or even 5-gallon containers. Although costs vary widely, weekly lab packs tend to cost \$1800 to \$3000 each week. When lab packs were conducted monthly, on the other hand, 55-gallon drums were routine. Records of monthly lab packs in 1996 and 1998 show costs of roughly \$5000 to \$7000 per month, about one-third less.³⁸

At least in principle, several methods exist to control flows. First, the quantities of chemicals used could be curtailed. Doing so would require less suitable experiments, fewer experimental trials, or greater efforts to prevent spills. Second, laboratories could manage their inventories more efficiently. Forgotten or misplaced stocks of unused chemicals linger on shelves until

³⁸The proper comparison, of course, should take the volume of waste into account. If generation has increased since the dates of those earlier lab packs, then the cost savings from monthly collection will be overstated. The high variability of flows, and the relatively short experience with weekly lab pack hauls, make a precise comparison difficult, but recent flows appear to be roughly comparable to those when lab packs took place monthly.

their expiration, while new chemical stocks are ordered. Separate storerooms for each lab lead to redundancy. At some cost, inventories could be combined, or managed to match use more closely. Finally, some fraction of the waste stream could be purified and reused rather than thrown away, albeit at significant cost.

Although some abatement would be desirable for Harvard’s chemistry labs, little concerted effort is made to curb flows. A partial explanation is that until recently, individual laboratories did not pay the costs of disposal, and thus had little incentive to cut back on chemical use. Individual labs recently began to pay a volume-based charge for both solvents and lab packs. Limited experience indicates that the use of chemical wastes – in particular, solvents – is fairly inelastic, suggesting high costs of substantial abatement.

5.2 Physical capital: investment and maintenance

Physical capital is subject to both stock and flow adjustments. Stock adjustments include adding to and replacing capital. Maintenance is the principal mechanism of flow adjustment; it slows depreciation. The optimal control of the quality of physical capital is central to a variety of problems, including the replacement of capital equipment, private investment decisions, and the purchase of durable goods. Treatments of these problems have focused on stock adjustment. Depreciation is either ignored, or treated as exogenous. In this section, we illustrate the parallels between a range of physical capital problems and the environmental quality problem considered above.

Consider the “machine-replacement problem”: a firm owns a piece of capital equipment, which deteriorates or becomes obsolete over time and periodically must be replaced. If replacement is the firm’s only option, this problem is equivalent to an optimal inventory problem, and an (S, s) -style rule is optimal. When the machine’s productivity falls to a certain point, perhaps measured in relation to the quality of a new machine, it is replaced. Rust (1987) and Cooper and Haltiwanger (1993) explore dynamic programming models of machine replacement that produce this result.³⁹ Two factors are implicit in their formulations. First, supplementing an old machine with a new one is not an option. Second, and more important

³⁹While both papers consider the machine replacement problem, neither dwells on it. Rust is concerned primarily with deriving an econometric specification to infer the parameters of utility and cost functions of managers from observed replacement patterns. Cooper and Haltiwanger are interested in exploring the aggregate implications of machine replacement across multiple firms. Neither paper incorporates maintenance as a choice variable in the formal analysis.

for our purposes, depreciation is exogenous – that is, maintenance is not considered.⁴⁰

General treatments of optimal investment under uncertainty also derive “trigger” rules similar to (S, s) policies. When the resale price (scrap value) of capital equipment is below its purchase price, and future demand for output is uncertain, optimal investment is determined by upper and lower bounds on the marginal revenue product of capital (Abel and Eberly 1994, 1996). The optimal rule is to purchase capital when the marginal revenue product exceeds the upper bound, and sell capital when the marginal revenue product falls below the lower bound. In the polar case of complete irreversibility, the lower bound is zero, and hence there is a single “trigger” above which investment is optimal (Pindyck 1988). Between the bounds, investment is zero, and depreciation proceeds apace.

The purchase and maintenance of consumer durables can be treated in similar fashion. Let wealth be variable, and assume a cost to adjust the stock of a durable, say a house. Consumers have a desired ratio of house value to wealth, possibly varying with wealth. If wealth moves up (or down) sufficiently, it becomes worthwhile to incur the transactions costs of selling one’s old house, and move into a new more (or less) expensive one. This problem is examined in detail by Grossman and Laroque (1990).

In all three of these problems, maintenance, which stems the expected downward flow in quality, may be subject to choice. Moreover, apart from the need to consider the capital’s scrap value, the costs of capital replacement and maintenance follow our assumptions above: the cost of a new piece of equipment is independent of the quality of the equipment it is replacing, and the costs of maintenance are likely to be convex in the rate of maintenance. Thus, the analysis above applies. Replacement and maintenance strategies should be jointly considered and jointly implemented.

The scrap value of physical capital will in general depend on quality. Our model above readily accommodates this consideration. The optimal value function, $J(x)$, would merely be computed with scrap value included for each value of x .

In section 5.4 below, we present a formulation that captures both the physical capital and human capital SFQ problems. It allows for both tech-

⁴⁰The optimal replacement of deteriorating assets – the general formulation of what we call the “machine replacement problem” – has also attracted considerable attention in the operations research literature. In the terms of our model, however, researchers in this literature (like Rust and Cooper and Haltiwanger in economics) have focused on restoration, rather than putting restoration and abatement together. For surveys, see Pierskalla and Voelker (1976) and Sherif and Smith (1981).

nological advance and scrap value.

5.3 Human capital: worker training

Investments in human capital – at least from the perspective of the firm – present an SFQ problem. Workers age, tire, and burn out. More important, in industries with rapid technological advance, workers' skills can quickly become obsolete. A firm can train its workers to maintain their productivity, but at some point it may lay off its older workers, or redeploy them to tasks where the latest technical skills are less essential, and replace them with recently trained workers. In the human capital context, training is the counterpart of curbing flows, while replacement amounts to restoration. The costs of these two activities must be balanced. Like environmental restoration, replacement entails significant fixed costs – severance payments, for example, or higher costs from experience-rated unemployment insurance.

Parallel to our model of pollution abatement, a firm's optimal investment in its workers is likely to vary over time, with reduced training as the workers near the trade-in point. Similarly, the literature on human capital has observed that a worker's expenditures on training should be high early in her career, since she will receive the payoffs from her investment over a longer period of time (Becker, 1964; Ben-Porath, 1967).

Interestingly, it may be the *relative* quality of the worker that deteriorates, rather than the absolute quality. (Similarly, technological obsolescence may be the force promoting the replacement of physical capital.) The downward drift in our model need not be a steady decline in a worker's actual productivity or a resource's quality; it may derive instead from a failure to keep up with a moving benchmark. If technological innovation raises the level of skills required by workers (as well as the level of skills acquired by newly trained ones), or diffusion of a new technology raises competitors' productivity, a firm faces an incentive to update or upgrade its capital just as surely as if the capital itself were deteriorating.

5.4 Optimal replacement and maintenance of physical or human capital: An example

We consider a physical capital example readily extrapolated to illustrate human capital as well. A firm must decide how much to spend to maintain its single machine and when to replace it. As a machine ages, its productivity tends to deteriorate: parts rust, gears wear, and so on. Despite this general downward trend, however, actual productivity may fluctuate with changes

in operating conditions, downtime, operating efficiency, the state of workers operating the machine, and so forth. Thus the productivity of an installed machine can be modeled as a Brownian motion with downward drift.

Similarly, the productivity of new-vintage machines can be modeled as a Brownian motion with positive drift. While we may assume that technology generally improves over time, there are two reasons why it may not follow an unimpeded upward path. First, the quality of craftsmanship may falter; a key producer may retire or leave the industry. Second, firms have heterogeneous needs. Workers differ from firm to firm, and the quality of other inputs into production may vary. As a result, some machines will work better than others for a given firm. Thus at the individual firm level, the productivity of the new vintage will have an upward drift, but may move down as well as up from one period to the next.

Consider a firm that produces output using a single machine at zero marginal cost, and sells output at a constant price of 1. At each time t , the firm operates a machine with an output rate y_t . The output rate of an installed machine evolves according to

$$dy_t = (q_t - \mu)dt + \sigma dw_{1,t},$$

where q_t is the maintenance rate at time t , μ is the average rate of deterioration, σ^2 is the variance rate of deterioration, and $w_{1,t}$ follows a standard Brownian motion. Maintaining the machine at a rate of q costs $c(q)$, a convex function.

At every point in time, the firm can choose to replace its installed machine with a new one. The output rate of a new-vintage machine purchased at time t is given by \bar{y}_t . This state of the art follows a Brownian motion $\bar{y}_t = \bar{y}_0 + \beta t - \rho w_{2,t}$, where $\beta > 0$ is the average rate of technological advance, ρ^2 is the variance rate, and $w_{2,t}$ is a standard Brownian motion that is independent from $w_{1,t}$. There is a fixed cost C of purchasing a new machine. The firm collects a scrap value of $S(y_t - \bar{y}_t)$ on the old machine; the further it is behind the current state of the art, the less it is worth.⁴¹

Letting τ_i be the i th time at which a new machine is purchased, the production rate for the firm evolves according to $y_t = y_0 + \int_{s=0}^t (q_s - \mu)ds + \sigma w_{1,t} - \sum_{\{i|\tau_i < t\}} (y_t - y_{\tau_i})$. The firm's objective is to maximize the present discounted value of expected profit, which equals the revenues from expected

⁴¹Given the constraint that each factory can only use one machine, relative productivity affects scrap value. An alternative but more complex formulation would have relative productivity matter because price declined with total output in the market.

production minus the costs of maintenance and replacement:

$$\max E \left[\int_{t=0}^{\infty} e^{-\alpha t} (y_t - c(q_t)) + \sum_{i=1}^{\infty} e^{-\alpha \tau_i} (S(y_t - \bar{y}_t) - C) \right],$$

where we have normalized the output price to 1. This problem can be formulated as a stochastic control problem with state variables q_t and \bar{q}_t , and the corresponding value function

$$J(y, \bar{y}) = \sup_{q_t, \tau_i} E_{y, \bar{y}} \left[\int_{t=0}^{\infty} e^{-\alpha t} (y_t - c(q_t)) + \sum_{i=1}^{\infty} e^{-\alpha \tau_i} (S(y_t - \bar{y}_t) - C) \right].$$

We can formulate the problem more succinctly in terms of the difference in quality between the currently operated machine and the state of the art. Let $x_t = y_t - \bar{y}_t$; then x_t is a Markov process evolving according to

$$(5) \quad x_t = x_0 + \int_{s=0}^t (q_s - \mu - \beta) ds + (\sigma^2 + \rho^2)^{1/2} w_t - \sum_{\{i | \tau_i < t\}} x_t$$

for some standard Brownian motion w_t . Note that x_t is negative when the old machine is inferior to the new vintage. Define a “differential value function”

$$(6) \quad V(x) = \sup_{q_t, \tau_i} E_{y, \bar{y}} \left[\int_{t=0}^{\infty} e^{-\alpha t} (x_t - c(q_t)) + \sum_{i=1}^{\infty} e^{-\alpha \tau_i} (S(x_t) - C) \right].$$

The difference between V and J does not depend on the state, since $V(x) = J(y, \bar{y}) + E \left[\int_{t=0}^{\infty} e^{-\alpha t} \bar{y}_t \right]$. Thus the differential value function V captures the relative quality of states as well as J does, and a pair (q_t, τ_i) attains the supremum for V if and only if it attains the supremum for J .

Applying the results of Section 4, the optimal policy will combine maintenance (abatement) with periodic replacement (restoration). We solve this problem numerically for particular parameter settings, provided in Table 1.

The optimal policy replaces the machine when its quality falls far enough below that of the new vintage. Between new machine purchases, the maintenance rate on the installed machine declines with its quality. Thus optimal maintenance on a newly installed machine is relatively high, but falls as the machine deteriorates and further still as replacement becomes imminent. Figure 6 illustrates the maintenance rate for the parameter values in Table 1. Note that the rate of repair does not fall all the way to zero, as it did in the earlier abatement case. This is a consequence of scrap value. Because old capital equipment has resale value which increases in its quality x ,

rate of technological advance	β	1
variance rate of technological advance	ρ^2	16
deterioration rate	μ	0.9
variance rate of deterioration	σ^2	4
discount rate	α	0.01
maintenance cost	$c(q)$	$20q^2$
maintenance ceiling	\bar{a}	20
cost of new machine	C	1000
scrap value	$S(x)$	$900 \min(e^{x/10}, 1)$

Table 1: Parameter settings for the capital replacement and maintenance problem.

maintenance always offers marginal benefits (even if only from increasing the scrap value).⁴² Moreover, for higher quality levels the rate of repair levels off at some constant, rather than decreasing in quality as it did earlier. This follows from the linearity of the profit function (as opposed to the strictly concave utility function used earlier).

The resulting output path for an individual firm is illustrated by Figure 7, which plots a simulation of the realized values for the paths of y_t and \bar{y}_t over time. Each machine replacement is marked by a vertical jump in the firm's output y_t . The difference in quality, x_t , is plotted near the bottom of the figure, and demonstrates the periodic replacements. Because depreciation and technical change are stochastic, the time between replacements varies widely. When deterioration is unusually rapid and the new-vintage machines improve dramatically, replacement is more frequent. When deterioration and obsolescence are more gradual, an installed machine is retained for a relatively long period before being scrapped and replaced.

The output of new-vintage machines, \bar{y}_t , drifts steadily upward over time. In general, the firm's output, y_t , lags behind, periodically catches up when the firm buys a new machine, and then lags behind again as the installed machine deteriorates and obsolesces. Occasionally, however, the output of the installed machine exceeds that of the new vintage. This might be the case, for example, when experience teaches the workers how better to operate the installed machine (i.e., learning by doing), while changes in the

⁴²More technically, to the right of the restoration trigger \underline{x} , the derivative of the value function J' is equal to that of scrap value, which is positive for all x . A smooth-pasting argument can be used to show that the derivative of the value function is continuous, and therefore the derivative cannot go to zero as x approaches the trigger. Thus the rate of repair does not go to zero either.




Figure 6: Optimal maintenance as a function of the difference in quality between installed and new vintage machines.




Figure 7: Simulated output paths of new vintage and installed machines, with the difference in output x_t .

new vintage prove to be of little use for the firm in question. In the event, the installed machine is retained for much longer than would otherwise be the case. Moreover, maintenance is higher over this period. As Figure 6 shows, maintenance is higher when the output of the installed machine is high relative to that of the new vintage.

The SFQ formulation is virtually the same if we consider a worker rather than a machine and assume, reasonably, that the cost of “trading her in,” e.g., severance or buyout payments, is inversely related to her quality (scrap value).

5.5 Multiple involved parties

Most of the examples above involve private decisions by firms. But SFQ models should also be understood by those who regulate systems. The government, for example, might be interested in promoting formal education and subsidizing training to bolster labor productivity. Policies in the two areas should be melded. When interventions affect stocks, or flows, but not both, indirect effects must also be considered. For physical capital, reducing the cost of restoration (e.g., through investment tax credits) will discourage maintenance, and vice versa.

The scope of government intervention varies across problems. In the case of pollution control, the pollution is a significant externality, but an arms-length approach will succeed. If the regulatory regime charges the externality cost to society on a continuous basis, e.g., through a quality-dependent pollution tax, the polluter will employ a hypothetical social planner’s optimal mix of restoration and abatement policies.

In many contexts where both restoration and abatement are possible, multiple decision makers may be operating with different incentives. A global optimization may be required to get the incentives and decisions right. For example, we examined the human capital problem from the standpoint of the firm. The firm views the worker to a substantial extent as an asset, whereas the worker views herself as the source of value. The employer will invest in training to the extent it staves off the various costs of switching to a new worker. But the worker who is forcibly retired might far prefer to keep working. Conceivably, workers’ concerns would be adequately recognized if the firm could communicate and commit to particular worker trade-in policies. If not, and if the worker reaps a rent from working at the firm, depreciated workers will be laid off too readily. Hence, wages will have to be higher to compensate, and the firm will suffer from its inability to commit.

The government, an involved third party through its tax collections, so-

cial welfare programs, and social security payments, and of course its concern for the welfare of its citizens, might prefer that the workers continue as well. If worker training is to be provided optimally, taking account of its effects on labor turnover and retirement, all involved parties should help pay for it. In practice, before the worker gets to the firm she and the government pay for training (schooling); once at the firm, training is financed by the firm and the government (through reduced taxes). School-leaving decisions and retirement decisions bring equivalent problems of nonalignment of payers and beneficiaries. SFQ models can identify the optimal mix of replacement and training for workers, and help design policies that better produce it. More modestly, they help to organize policy discussion when the maintenance of quality is a concern.

6 Conclusion

In a wide range of settings, both stocks and flows can be controlled to affect the quality of a resource – whether a natural resource, physical capital, or human capital. In these cases, the SFQ model applies. Managing the resource entails curbing the downward drift in quality and the periodic replenishment of the stock.

Though our model can address physical and human capital, both subject to depreciation, we focus primarily on the optimal management of an environmental resource subject to continuing degradation. Cost structures play a key role in our analysis. Abatement is taken to follow conventional convexity assumptions, so that the cost of abatement depends on the magnitude of stanching flow. Restoration, by contrast, is assumed to offer significant economies of scale in the magnitude of cleanup.

We focus on the intertwined roles of abatement and restoration, analyzing the limiting cases when only one of these methods is available to provide context for our results. When only abatement is possible, the optimal policy (as has been well documented in the literature) is to progress to a steady-state, or (given stochastic processes) to what we call an “expectation equilibrium,” where abatement efforts just offset expected damage to the resource. When only restoration is possible, the optimal policy establishes a “trigger level” of environmental quality. Damage accumulates and the quality of the resource declines until the trigger level is reached; the resource is restored to a high level of quality, and the process repeats.

The most interesting case arises when both restoration and abatement are available and the full SFQ model comes into play. The optimal policy

calibrates the use of these two instruments. The greater the flow, and the lower the restoration cost, the more restoration is relied upon. Abatement and restoration are not separable: that is, the feasibility of restoration affects the optimal abatement path, and the feasibility of abatement influences the nature of restoration decisions.

We have also shown that the principles emerging from the management of natural resources offer insights into the optimal management of both physical and human capital. The balance between maintaining a piece of capital equipment and replacing it, or between training a worker and replacing her, follow much the same logic as harmonizing abatement of pollution flows and restoration of a resource.

PhD. Candidate in Political Economy and Government, Harvard University; nkeohane@fas.harvard.edu;

Departments of Management Science and Engineering and Electrical Engineering, Stanford University; bvr@stanford.edu;

John F. Kennedy School of Government, Harvard University, and NBER; richard_zeckhauser@harvard.edu.

REFERENCES

- Abel, A. B., and J. C. Eberly (1994): "A Unified Model of Investment Under Uncertainty," *American Economic Review*, 84, 1369-1384.
- (1996): "Optimal Investment with Costly Reversibility," *Review of Economic Studies*, 63, 581-593.
- Arrow, K. J., T. Harris, and J. Marschak (1951): "Optimal Inventory Policy," *Econometrica*, 19, 250-272.
- Balcer, Y. M. (1973): "Optimal Advertising and Inventory Policy with Random Demand." Doctoral dissertation, Department of Operations Research, Stanford University.
- Becker, G. S. (1975): *Human Capital*, 2nd ed. Chicago: University of Chicago Press.
- Ben-Porath, Y. (1967): "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 75, 352-365.
- Caputo, M. R., and J. E. Wilen (1995): "Optimal Cleanup of Hazardous Wastes," *International Economic Review*, 36, 217-243.
- Cooper, R., and J. Haltiwanger (1993): "Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, 83, 360-382.
- Dasgupta, P. S., and G. M. Heal (1979): *Economic Theory and Exhaustible Resources*. Cambridge, UK: Cambridge University Press.
- Falk, I., and R. Mendelsohn (1993): "The Economics of Controlling Stock Pollutants: An Efficient Strategy for Greenhouse Gases," *Journal of Environmental Economics and Management*, 25, 76-88.
- Grossman, S., and G. Laroque (1990): "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58, 25-51.
- Hotelling, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39, 137-175.
- Keeler, E., M. Spence, and R. Zeckhauser (1971): "The Optimal Control of Pollution," *Journal of Economic Theory*, 4, 19-34.

- Krylov, N. V. (1980): *Controlled Diffusion Processes*. New York: Springer-Verlag.
- Kushner, H. J., and P. G. Dupuis (1992): *Numerical Methods for Stochastic Control Problems in Continuous Time*. New York: Springer-Verlag.
- Phillips, C. V., and R. J. Zeckhauser (1998): "Restoring Natural Resources With Destination-Driven Costs," *Journal of Environmental Economics and Management*, 36, 225-42.
- Pierskalla, W. P., and J. A. Voelker (1976): "A Survey of Maintenance Models: The Control and Surveillance of Deteriorating Systems," *Naval Research Logistics Quarterly*, 23, 353-388.
- Pindyck, R. S. (1988): "Irreversible Investment, Capacity Choice, and the Value of the Firm," *American Economic Review*, 78, 969-985.
- Plourde, C. G. (1972): "A Model of Waste Accumulation and Disposal," *Canadian Journal of Economics*, 5, 119-125.
- Plourde, C., and D. Yeung (1989): "A Model of Industrial Pollution in a Stochastic Environment," *Journal of Environmental Economics and Management*, 16, 97-105.
- Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999-1033.
- Scarf, H. (1960): "The Optimality of (S,s) Policies for the Dynamic Inventory Problem," *Proceedings of the First Stanford Symposium on Mathematical Methods in Social Sciences*. Stanford, CA: Stanford University Press.
- Sherif, Y.S., and M. L. Smith (1981), "Optimal Maintenance Models for Systems Subject to Failure – A Review," *Naval Research Logistics Quarterly*, 32, 47-74.
- Smith, V. L. (1972): "Dynamics of Waste Accumulation: Disposal Versus Recycling," *Quarterly Journal of Economics*, 86, 600-616.
- Stokey, N. L. and R. E. Lucas, Jr. (1989): *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.

APPENDIX A: Proofs of Theorems

Proof of Theorem 1

We have $J < 0$ because $u < 0$ and $c > 0$. Furthermore, for each x , $J(x) > -\infty$ because $E_x [\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt]$, the expected utility in the absence of abatement, is finite. We have verified Assertion 1.

Given the finiteness of J , it follows from standard results in the theory of controlled diffusions (see, e.g., Theorem 5 on pages 24–25 of (Krylov, 1980)) that J is twice continuously differentiable, and that

$$\sup_{q \in [0, \bar{a}]} \left(\frac{\sigma^2}{2} J''(x) + (q - \mu) J'(x) - \alpha J(x) + u(x) - c(q) \right) = 0$$

for all x , rendering Assertions 2 and 3.

Using straightforward sample–path arguments, it is easy to establish that J is increasing and strictly concave and that J' is unbounded above, as consequences of similar properties enjoyed by u . Assertions 4, 5, and 6, follow.

Let a function f_x be defined by $f_x(q) = qJ'(x) - c(q)$. Note that $f'_x(0) = J'(x) > 0$ and $f''_x = -c'' \leq -\epsilon$ for some ϵ . Hence, for any x , the supremum

$$(7) \quad \sup_{q \in [0, \bar{a}]} f_x(q)$$

is uniquely attained by some $a^*(x) \in [0, \bar{a}]$. It is easy to see that $a^*(x)$ also attains the supremum in the Hamilton–Jacobi–Bellman equation. The fact that a^* is an optimal policy follows from standard results in the theory of controlled diffusions. We thereby have Assertions 7 and 8.

Q.E.D.

Proof of Theorem 2

Once again, let f_x be defined by $f_x(q) = qJ'(x) - c(q)$, so that $a^*(x)$ uniquely attains the supremum in $[0, \bar{a}]$ of f_x . (Recall that $f'_x(0) = J'(x) > 0$ and $f''_x = -c'' \leq -\epsilon$ for some ϵ .) Consider the less constrained problem

$$(8) \quad \sup_{z \in [0, \infty)} f_x(z).$$

Since $f'' \leq -\epsilon$, the supremum is always attained by some $z \in (0, \infty)$. Let $b(x)$ denote the optimum for a given state x . Because $f'_x(0) > 0$, $b(x) > 0$. Furthermore, since $f'_x(z)$ decreases as x increases, b is decreasing.

It is easy to see that $a^*(x) = \min(b(x), \bar{a})$. Since J' is unbounded below, for any $z > 0$ there exists a state x such that $f'_x(z) > 0$, implying that b is unbounded above, and therefore, there exists a state \hat{x} such that $a^*(x) = \bar{a}$ for $x \leq \hat{x}$. Assertion 1 follows.

Recall that $J < 0$ and $J' > 0$, so that $\lim_{x \rightarrow \infty} J'(x) = 0$. Hence for any $z > 0$, there exists a state x such that $f'_x(z) < 0$, implying that $\lim_{x \rightarrow \infty} b(x) = 0$ and that Assertion 2 holds. The fact that b is decreasing implies that there exists a state x^* such that $\mu < b(x)$ for $x < x^*$ and $\mu > b(x)$ for $x > x^*$. Since $\mu < \bar{a}$, we have Assertion 3.

Q.E.D.

Proof of Theorem 3

Since the restoration-only case is a special case of the abatement-and-restoration case, with the abatement constraint $\bar{a} = 0$, this theorem is a corollary to Theorem 5. We therefore omit the proof.

Proof of Theorem 4

Let $J(x, C)$ denote the optimal value of state x given a restoration cost $C > 0$. It is easy to show by a sample path argument that for any x , $J(x, C)$ is decreasing in C . Fix $C_2 > C_1 > 0$ and assume for contradiction that $\underline{x}(C_2) \geq \underline{x}(C_1)$. Let $\tau = \inf\{t | x_t = \underline{x}(C_2)\}$. We then have

$$\begin{aligned} J(0, C_2) &= E_0 \left[\int_{t=0}^{\tau} e^{-\alpha t} u(x_t) dt + e^{-\alpha \tau} J(\underline{x}(C_2), C_2) \right] \\ &= E_0 \left[\int_{t=0}^{\tau} e^{-\alpha t} u(x_t) dt + e^{-\alpha \tau} J(\underline{x}(C_2), C_1) \right] \\ &\quad + E_0 [e^{-\alpha \tau} (J(\underline{x}(C_2), C_2) - J(\underline{x}(C_2), C_1))] \\ &= J(0, C_1) + E_0 [e^{-\alpha \tau} (J(\underline{x}(C_2), C_2) - J(\underline{x}(C_2), C_1))] \\ &< J(0, C_1) + J(\underline{x}(C_2), C_2) - J(\underline{x}(C_2), C_1). \end{aligned}$$

It follows that

$$J(0, C_2) - J(\underline{x}(C_2), C_2) < J(0, C_1) - J(\underline{x}(C_2), C_1) \leq J(0, C_1) - J(\underline{x}(C_1), C_1), \quad (9)$$

where the final inequality relies on our assumption that $\underline{x}(C_2) \geq \underline{x}(C_1)$.

Theorem 3 asserts that for any $C > 0$, $J(0, C) - J(\underline{x}(C), C) = C$. Inequality 9 therefore implies that $C_2 < C_1$, which yields a contradiction.

Q.E.D.

Proof of Theorem 5

We have $J < 0$ because $u < 0$, $c > 0$, and $C > 0$. Furthermore, for each x , $J(x) > -\infty$ because $E_x [\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt]$ is finite. We have established Assertion 1.

Because restoration sets the state to 0 and costs C , $J(x) \geq J(0) - C$ for all x , and an optimal policy R can be defined to be the set of all x such that $J(x) = J(0) - C$. Let us establish that any optimal policy R is nonempty – that is, that at some level of environmental quality the manager restores the resource. Assume, for contradiction, that the optimal restoration policy R is empty. Then, we would have $J(x) = \sup_a E_x^a [\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt]$, which corresponds to the optimal value function for the case where only abatement is possible. By Theorem 1, J would be unbounded below, contradicting the fact that $J(x) \geq J(0) - C$.

By straightforward sample-path arguments, it is easy to show that J is continuous and nondecreasing. Hence, there exists a state \underline{x} such that $J(x) = J(0) - C$ for all $x \leq \underline{x}$ and $J(x) > J(0) - C$ for all $x > \underline{x}$, rendering Assertion 2. It follows that the restoration component of an optimal policy is given by $R^* = (-\infty, \underline{x}]$.

It follows from Theorem 3 on page 39 of Krylov (1980) that J is twice continuously differentiable on (\underline{x}, ∞) and differentiable everywhere. Furthermore, J satisfies

$$\sup_{q \in [0, \bar{a}]} \left(\frac{\sigma^2}{2} J''(x) + (q - \mu) J'(x) - \alpha J(x) + u(x) - c(q) \right) = 0$$

for all $x > \underline{x}$. Hence, Assertions 3, 4, and 5 are valid. It is easily verified by sample-path arguments that J is increasing on (\underline{x}, ∞) , and Assertions 6 and 7 follow.

It follows from Assertions 2 and 4 that $J'(\underline{x}) = 0$. Since $J'(x) > 0$ for all $x > \underline{x}$, we have $J''(x) > 0$ on some range $x \in (\underline{x}, y)$ for some $y > \underline{x}$. Furthermore, since J is bounded above, $J''(x)$ must be negative for some $x > \underline{x}$, and by continuity of the second derivative, there is a well-defined minimal inflection point $x^\dagger = \min\{x > \underline{x} | J''(x) = 0\}$, which by definition satisfies Assertions 8 and 9.

Let a function f_x be defined for $x > \underline{x}$ by $f_x(q) = qJ'(x) - c(q)$. Note that $f_x'' = -c'' \leq -\epsilon$ for some ϵ . Hence, for any x , the supremum

$$(10) \quad \sup_{q \in [0, \bar{a}]} f_x(q)$$

is uniquely attained by some $q \in [0, \bar{a}]$. For each state $x > \underline{x}$, let $a^*(x)$ be the value attaining the supremum, and note that (a^*, R^*) constitutes

an optimal policy since the values $a^*(x)$ also attain the supremum in the Hamilton–Jacobi–Bellman equation (equation 4). This validates Assertions 11 and 12.

We are left with the task of establishing Assertion 10. Given scalars $\Delta > 0$ and $x > x^\dagger + \Delta$, we define two processes

$$x_t^- = x + \int_{s=0}^t (a^*(x_s^-) - \mu) dt + \sigma w_t,$$

and

$$x_t^+ = x + 2\Delta + \int_{s=0}^t (a^*(x_s^+) - \mu) dt + \sigma w_t,$$

each evolving on $[0, \tau]$, where τ is given by

$$\tau = \inf\{t | x_t^- = x^\dagger \text{ or } x_t^- = x_t^+\}.$$

Let

$$x_t = x + \Delta + \int_{s=0}^t \left((a^*(x_s^+) + a^*(x_s^-))/2 - \mu \right) dt + \sigma w_t,$$

and note that $x_t = (x_t^+ + x_t^-)/2$ for all $t \in [0, \tau]$. It is easy to show that τ is finite with probability one.

Define “sample costs” associated with the three processes:

$$\hat{J}(x, \omega) = \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t) - c((a^*(x_t^+) + a^*(x_t^-))/2) \right) dt + e^{-\alpha \tau} J(x_\tau),$$

$$\hat{J}^+(x, \omega) = \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t^+) - c(a^*(x_t^+)) \right) dt + e^{-\alpha \tau} J(x_\tau^+),$$

$$\hat{J}^-(x, \omega) = \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t^-) - c(a^*(x_t^-)) \right) dt + e^{-\alpha \tau} J(x_\tau^-),$$

where ω denotes the sample path of the underlying Brownian motion w_t .

We will show that for almost all ω and any $x \in (x^\dagger, \infty)$,

$$\hat{J}(x, \omega) \geq \frac{1}{2} \left(\hat{J}^+(x, \omega) + \hat{J}^-(x, \omega) \right).$$

We consider two separate cases that together comprise a set of probability 1. The first is when $x_\tau^- \neq x^\dagger$. In this event, we have $x_\tau^- = x_\tau^+ = x_\tau > x^\dagger$, and the desired inequality follows directly from concavity of u and convexity of c .

The second case is when $x_\tau^- = x^\dagger$. Given our assumptions on c , the fact that a^* is bounded above, and the fact that J is bounded and twice

continuously differentiable on (\underline{x}, ∞) , it can be shown that for any $y > \underline{x}$, $|a^*(y) - a^*(y + \Delta)| = O(\Delta)$. It follows that $\sup_{t \in [0, \tau]} |a^*(x_t^-) - a^*(x_t^+)| = O(\Delta)$ and $x_\tau^+ - x^\dagger = O(\Delta)$. We then have

$$\begin{aligned}
\hat{J}(x, \omega) &= \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t) - c((a^*(x_t^+) + a^*(x_t^-))/2) \right) dt + e^{-\alpha \tau} J(x_\tau) \\
&\geq \frac{1}{2} \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t^+) + u(x_t^-) - c(a^*(x_t^+)) - c(a^*(x_t^-)) \right) dt + e^{-\alpha \tau} J(x_\tau) \\
&= \frac{1}{2} \int_{t=0}^{\tau} e^{-\alpha t} \left(u(x_t^+) + u(x_t^-) - c(a^*(x_t^+)) - c(a^*(x_t^-)) \right) dt \\
&\quad + \frac{1}{2} e^{-\alpha \tau} \left(J(x_\tau^+) + J(x^\dagger) + o(\Delta^2) \right) \\
&= \frac{1}{2} \left(\hat{J}^+(x, \omega) + \hat{J}^-(x, \omega) + o(\Delta^2) \right),
\end{aligned}$$

where the second-to-last expression relies on the fact that $J''(x^\dagger) = 0$ and that $x_\tau^+ - x^\dagger = O(\Delta)$. It follows that for almost all ω and any $x \in (x^\dagger, \infty)$, $\hat{J}(x, \omega)$ is concave in x .

By Bellman's principle of optimality, we have

$$J(x) = E[\hat{J}^-(x, \omega)], \quad J(x + 2\Delta) = E[\hat{J}^+(x, \omega)], \quad \text{and} \quad J(x + \Delta) \geq E[\hat{J}(x, \omega)].$$

Hence,

$$J(x + \Delta) \geq \frac{1}{2} \left(J(x) + J(x + 2\Delta) + o(\Delta^2) \right),$$

and therefore $J''(x) < 0$ for $x > x^\dagger$.

Q.E.D.

Proof of Theorem 6

As a step toward establishing Assertion 1, we will show that $\tilde{J} < J$. It is easy to see that $\tilde{J} \leq J$. From Theorem 5, we have $J'(\underline{x}) = 0 < \tilde{J}'(\underline{x})$. This implies that $\tilde{J}(\underline{x}) < J(\underline{x})$. For $x < \underline{x}$, we then have $J(x) = J(\underline{x}) > \tilde{J}(\underline{x}) > \tilde{J}(x)$. For $x > \underline{x}$, on the other hand, the fact that $\tilde{J}(x) < J(x)$ follows from our observation that $\tilde{J}(\underline{x}) < J(\underline{x})$ coupled with standard sample-path arguments.

Consider two states y and z with $\underline{x} \leq y < z$. By Bellman's principle of optimality (see, e.g., Krylov), we have

$$\tilde{J}(z) = \sup_a E_z^a \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt + e^{-\alpha \tau} \tilde{J}(y) \right]$$

and

$$\begin{aligned} J(z) &= \sup_{a, R} E_z^{a, R} \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt + e^{-\alpha \tau} J(y) \right] \\ &= \sup_a E_z^a \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt + e^{-\alpha \tau} J(y) \right], \end{aligned}$$

where τ is the first time at which $x_t = y$. (The final equality holds because $x_t > \underline{x}$ for $t \leq \tau$.)

Let \tilde{a} be an optimal policy for the case where only abatement is possible. We then have

$$\begin{aligned} J(z) - \tilde{J}(z) &= \sup_a E_z^a \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt + e^{-\alpha \tau} J(y) \right] \\ &\quad - \sup_a E_z^a \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt + e^{-\alpha \tau} \tilde{J}(y) \right] \\ &\leq E_z^{a^*} \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a^*(x_t))) dt + e^{-\alpha \tau} J(y) \right] \\ &\quad - E_z^{a^*} \left[\int_{t=0}^{\tau} e^{-\alpha t} (u(x_t) - c(a^*(x_t))) dt + e^{-\alpha \tau} \tilde{J}(y) \right] \\ &= E_z^{a^*} \left[e^{-\alpha \tau} (J(y) - \tilde{J}(y)) \right] \\ &< J(y) - \tilde{J}(y). \end{aligned}$$

It follows that $J' < \tilde{J}'$, which gives us Assertion 1.

We now turn to Assertion 2. Once again, let a function f_x be defined for $x > \underline{x}$ by $f_x(q) = qJ'(x) - c(q)$. Recall that $f_x'' = -c'' \leq -\epsilon$ for some ϵ and that for each $x > \underline{x}$, $a^*(x)$ uniquely attains the supremum of f_x , as well as that of the Hamilton–Jacobi–Bellman equation. For any $x, y \in (\underline{x}, x^\dagger)$ with $x < y$, $f_y'(a^*(x)) > f_x'(a^*(x)) = 0$, since $J'' > 0$ on $(\underline{x}, x^\dagger)$. Consequently, unless $a^*(x) = \bar{a}$, we have $a^*(y) > a^*(x)$. An entirely analogous argument establishes that $a^*(y) < a^*(x)$ if $x^\dagger < x < y$ and $a^*(y) \neq \bar{a}$. Assertion 2 follows.

Finally, consider Assertion 3. Let \tilde{f}_x be defined by

$$\tilde{f}_x(q) = q\tilde{J}'(x) - c(q).$$

Recall that for any x , the supremum of \tilde{f}_x is uniquely attained by $\tilde{a}(x)$. Since $\tilde{J}' > J'$, for every $x > \underline{x}$, we have $\tilde{f}_x'(a^*(x)) > f_x'(a^*(x))$. This implies that if $a^*(x) < \bar{a}$ then $\tilde{a}(x) > a^*(x)$. Hence, we have Assertion 3.

Q.E.D.

variance	σ^2	9
discount rate	α	0.005
utility	$u(x)$	$-e^{-0.05(x+150)}$
abatement cost	$c(q)$	$40q^2$
abatement ceiling	\bar{a}	20
restoration cost	C	13000

Table 2: Parameter settings for computations in Section 4.

APPENDIX B: Computations of Section 4.3

Computations that generated the plots appearing in Section 4.3 were conducted using specific functional forms and fixed parameter values. The exception was the flow rate μ , which varied as indicated in the figures. The functional forms were a quadratic function for the abatement cost and a negative natural exponential function for utility. The fixed parameter values and exact functional forms are provided in Table 2.

Value functions were computed via policy iteration on a “locally consistent” approximating Markov chain (see, e.g., Kushner and Dupuis, 1992).