The Optimal Harvesting of Environmental Bads

Nathaniel Keohane
Harvard University
nkeohane@fas.harvard.edu

Benjamin Van Roy
Stanford University
bvr@stanford.edu

Richard Zeckhauser
Harvard University
richard_zeckhauser@Harvard.Edu

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.

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. 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 management 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 present a general model of the optimal management of natural resources when restoration (with economies of scale) is an option. 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. Due to space limitations, we provide only a brief account of some central results. For a more comprehensive exposition, including an overview of previous literature, proofs, and applications of the ideas to environmental policy and management of human and physical capital, we refer the reader to our full–length paper (Keohane et al, 2000).

2 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. 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.

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 $z_0 = 0$. Hence, damages evolve according to

$$z_t = \mu t - \sigma w_t, \tag{1}$$

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

2.1 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.

The resource manager's problem is to determine the optimal abatement rate.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). 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 $a : \Re \mapsto [0, \overline{a}]$ 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.Under an abatement policy a, the state of 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 $\epsilon > 0$.*

2.2 Restoration Policies

Next, consider a setting in which 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.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 $x = 0$. Letting τ_i denote the ith 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.

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 nonconvexity, 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.

A restoration policy is characterized by a measurable closed subset R of $\Re$. 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.

2.3 Combined Abatement–Restoration Policies

Our focus is on 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.

3 Optimal policies for the SFQ problem

As we will show in this section, when both abatement and restoration are possible, both are employed. Moreover, the availability of restoration affects the optimal abatement policy, and the possibility of abatement alters restoration. Theorems 1 and 2 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 1 *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 $x \le \underline{x}$;*
3. *J is twice continuously differentiable on $(\underline{x}, \infty)$;*
4. *J is continuously differentiable;*
5. *J satisfies*

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

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

6. *$J' \ge 0$;*
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 $a^* : (\underline{x}, \infty) \mapsto [0, \overline{a}]$ such that for every $x \in (\underline{x}, \infty)$, $a^*(x)$ uniquely attains the supremum in equation (2); 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.

The next theorem describes the optimal combined abatement-restoration policy (a^*, R^*) defined by Theorem 1, 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.

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

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

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. 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. Note that when restoration is 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

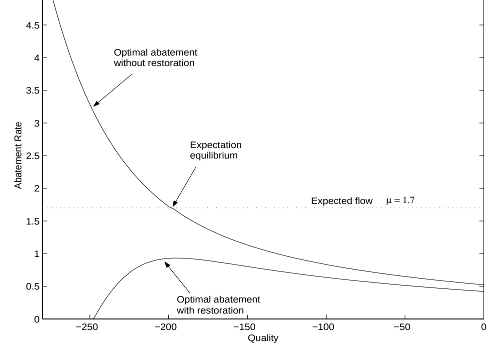

Figure 1: Optimal abatement policies with and without the possibility of restoration.

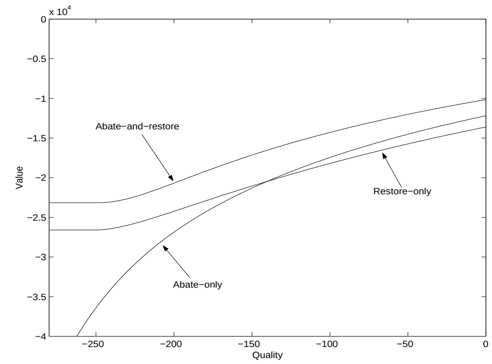

Figure 2: Value functions for the abate–only, restore–only, and abate–and–restore cases.

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 $\underline{x}$.

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 restoration 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^{**}. The higher value, x^*, is a stable expectation equilibrium. Once the quality of the resource hits x^*, it will tend to

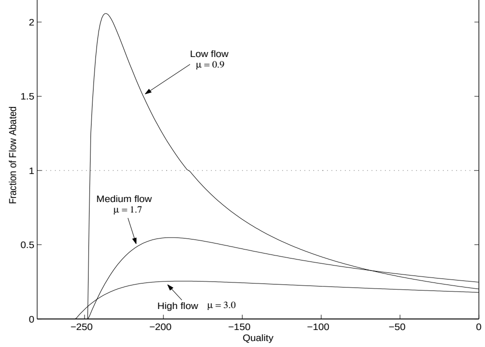

Figure 3: Fraction of flow abated as a function of quality, for three flow rates.

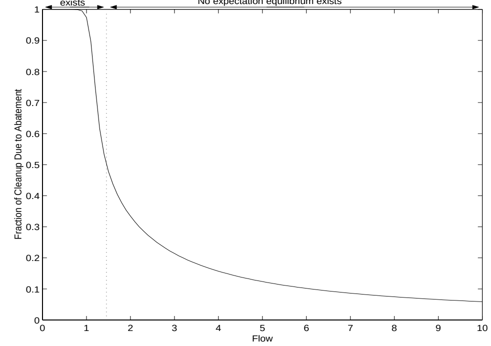

Figure 4: Fraction of cleanup due to abatement, as a function of the mean flow rate μ.

oscillate around that level for small perturbations in the flow of damages. Small stochastic increases in the flow above 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. 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 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 restora-

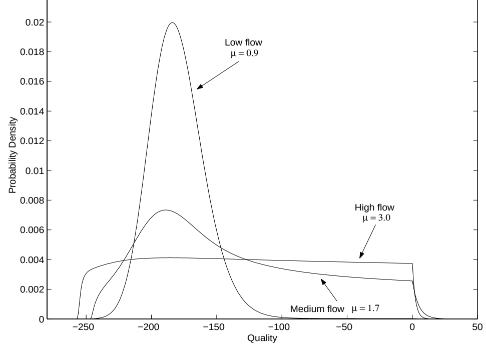

Figure 5: Frequency distributions of resource qualities (states) under optimal policies for three flow rates.

tions 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. 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 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.

REFERENCES

Keohane, N. O., B. Van Roy, and R. Zeckhauser (2000): "Controlling Stocks and Flows to Promote Quality: The Environment, with Applications to Physical and Human Capital," NBER working paper No. 7727.

Krylov, N. V. (1980): *Controlled Diffusion Processes.* New York: Springer–Verlag.