

Managing the Quality of a Resource with Stock and Flow Controls*

Nathaniel O. Keohane [†] Benjamin Van Roy Richard J. Zeckhauser
Yale University Stanford University Harvard University

June 8, 2006

Abstract

We consider a class of problems, which we call “SFQ” problems, in which both stocks and flows can be controlled to promote the quality of a valued resource, such as environmental quality or public infrastructure. Under the optimal policy, periodic restoration of the stock of quality complements positive but variable abatement of the flow of damages. When deterioration is more rapid or highly variable, or when abatement is more expensive relative to restoration, the optimal policy relies relatively more on restoration.

When deterioration is due to private firms or individuals, a flow tax equal to the present value of marginal damages provides efficient incentives for abatement. This tax rises at first as quality worsens, but eventually falls as restoration nears. The revenues raised by such a tax approximates the cost of restoration, with the two quantities converging as the variance of flows goes to zero.

We discuss the implications of the SFQ model for a range of real-world problems in the environmental arena, and for the management of public infrastructure. But the lessons are general, and we briefly discuss how they apply to private stocks of physical and human capital.

JEL Nos. C61, H54, Q20, Q50.

Keywords: abatement, restoration, pollution, public infrastructure, (S,s) policies, stochastic control theory.

*The authors thank, without implicating, Miriam Avins, Bill Brainard, George Borjas, David Ellwood, Eduardo Engel, Carolyn Fischer, Peter Glynn, George Hall, Sandy Jencks, David Laibson, John Leahy, Barry Nalebuff, Ben Polak, Dan Schrag, Pete Veinott, two anonymous referees, and seminar participants at Stanford University, Harvard University, and the NBER Summer Institute. Tyler Welti performed excellent research assistance.

[†]Corresponding author. Mailing address: Nathaniel Keohane, Yale School of Management, 135 Prospect St., P.O. Box 208200, New Haven, CT 06511; email nathaniel.keohane@yale.edu; telephone (203) 432-6024; fax (203) 432-6974.

1 Introduction

Public policies across a range of issues involve maintaining the quality of a valued resource stock. Thus, we try to prevent the deterioration of environmental quality and to keep our roads well paved. The economist's typical prescription in such settings is to identify the level where the marginal cost of maintaining quality equals its marginal benefit, and then to stay there. Policy proposals for global warming, for example, seek to identify a target concentration of carbon dioxide in the atmosphere. Such an approach is optimal when quality can be maintained only by slowing the rate of deterioration, e.g., through pollution abatement. However, there are often economies of scale in moving a stock to a desired level of quality. In such cases, periodic restoration of the stock will be worthwhile, and the optimal rate of abatement will adjust.

A simple example illustrates our theory. Consider the Longfellow Bridge linking Cambridge and Boston. The rate at which it deteriorates can be slowed by regular maintenance, or by traffic restrictions such as weight limits. Eventually, however, it will wear out and have to be replaced. (It was completed a century ago, then replacing a century-old bridge.) Over the course of its lifetime, optimal maintenance for the bridge will vary. Early on, maintenance increases with time, as it requires increasing attention. Toward the end of its life, however, optimal maintenance declines, since the future benefits of better quality for the current bridge decline. There is no need to fill the cracks in a bridge that will soon be replaced.

The motivating observation of this paper is that the simple logic of the maintenance and replacement of a bridge applies to the management of the quality of a wide range of valued resources. A canonical example in the environmental arena is the accumulation and treatment of waste at landfills or generating sites. Optimal management both slows the generation of new wastes is slowed and periodically cleans up accumulated stocks. The latter approach – capping a landfill or clearing out a storage area – capitalizes on scale economies. Similar dynamics play out in a variety of other instances: a reservoir fills with sediment until the dam must be replaced; a groundwater aquifer is drawn down and eventually recharged; potholes are patched until a street is repaved; drug dealing is tamped down by a police presence until a massive drug sweep is undertaken.

In each of these settings, two distinct approaches are available to manage the resource: boosting its quality, and slowing the rate at which it deteriorates. Hence both *stocks* and *flows* can be controlled to promote *quality*. We use the acronym “SFQ” to refer to this class of problems. In this paper, we develop a general model of the optimal management of a resource stock when flows

are controllable and restoration of the stock is feasible. We discuss how the model applies to environmental quality and public infrastructure. The lessons are general, however, and extend readily to a variety of private applications, such as physical or human capital stocks within a firm.

We assume that the costs of flow control (which we call *abatement*) increase on the margin, but that stock control (or *restoration*) exhibits economies of scale, so that discrete improvements are potentially desirable. Such scale economies are likely to obtain in many settings. For example, cleaning up a hazardous waste site typically requires hauling the soil away for off-site incineration, in which case the costs vary little with the concentration of the contaminant in the soil. Similarly, there are high fixed costs involved in dredging a river or paving city streets. As will become apparent below, what is crucial to our analysis is that there are economies of scale “at the bottom” – that is, that the costs of restoration do not increase too rapidly as the quality of the stock diminishes. In many settings, the source of the nonconvexity is institutional rather than technological. Consider fiscal policy. Over time, the tax code (the “stock” in this example) accumulates loopholes and exceptions that distort incentives and undermine efficiency. While tax simplification requires substantial political capital, observation suggests that the political costs are little different for small and large reforms. Thus, tax reform is a rare and lumpy event, eliminating many loopholes at once.

Given economies of scale in restoration, the optimal policy calls for restoring the resource whenever quality falls to a sufficiently low level. At states above that point, the new flow is abated at a rate that varies with the current quality of the resource. After restoration occurs, deterioration resumes, quality starts to decline (albeit stochastically), and the cycle repeats. The optimal trade-off between abatement and restoration depends on the magnitude and variability of flows, the relative costs of the two strategies, and the discount rate. If flows are low enough, or if abatement is cheap enough, optimal abatement may rise to offset expected deterioration, achieving an equilibrium in expectation. Even in this case, restoration will occur if unexpected shocks reduce the stock of quality sufficiently; hence its availability influences the optimal abatement path. When deterioration is more rapid or more variable, or when restoration is relatively less costly, the optimal policy relies more on restoration.

In many settings, the stock deteriorates at a speed determined by the actions of myriad firms or individuals. The central authority must now align private incentives with social welfare. We consider the natural case in which abatement must be undertaken by a large number of firms but restoration is implemented by the center. A municipal landfill is a natural example. Private parties (perhaps influenced by taxes or other policy measures) control the flow into it. Ultimately,

the government caps it and restores the site. In such a case, the optimal abatement path can be achieved by charging a time-varying flow tax equal to the present value of marginal damages. The tax has a surprising pattern. As the quality of the stock decreases, this optimal tax rises at first, but it eventually falls as the state worsens and restoration nears. Moreover, the optimal tax rate may be lower when there is more pressure on quality (e.g., greater unregulated waste flow). We also consider the tax as a source of revenue. Raising funds to pay for restoration might appear to be independent from aligning private incentives. We show that revenues from the flow tax approximate the cost of restoration, with the two quantities converging as the variance of flows goes to zero.

Our analysis melds two instruments that have typically been considered in isolation. The theoretical literature on capital investment has centered on replacement (restoration) rather than maintenance (abatement).¹ Optimal investment in these models typically follows an (S,s) policy (Arrow, Harris, and Marschak 1951). On the other hand, conventional models of the optimal management of stock pollutants have modeled abatement alone. The optimal policy in that setting equates the marginal benefit of reducing pollution, adjusted for the discount rate and the decay rate of the stock, to the marginal cost of abating it. A steady state is reached in which optimal abatement efforts just keep up with net new accumulation (Falk and Mendelsohn 1993; Keeler, Spence, and Zeckhauser 1971; Plourde 1972; Plourde and Yeung 1989; Smith 1972).² During the transition to the steady state, the shadow value of environmental quality rises steadily, and the optimal tax rises with it (Farzin 1996). In contrast, when restoration offering economies of scale is available, the optimal policy may entail periodic restorations punctuating long periods of deterioration that are only partially offset by abatement, with the optimal tax rate first rising and then falling as resource quality worsens.

The next section introduces the formal model. Section 3 characterizes optimal abatement and restoration policies, and Section 4 discusses the optimal tax in a decentralized setting. Section 5 illustrates our analysis with real-world examples. Section 6 concludes.

¹For models of investment in physical capital, see Feldstein and Rothschild (1974) and Abel and Eberly (1994). Nickell (1975) considers maintenance as well as replacement, but models maintenance as exogenously determined. For a model of optimal consumption of durable goods, see Grossman and Laroque (1990).

²A few models of the optimal cleanup of an accumulated pollution stock have considered restoration but not abatement. Caputo and Wilen (1995) assume that cleanup costs are convex. As a result, the optimal solution stops short of complete cleanup (they let natural degradation finish the process), as long as when pollution approaches zero so does its marginal damage. Phillips and Zeckhauser (1998) assume economies of scale in cleanup, but consider the problem in a static setting and hence ignore abatement.

2 Model framework

Our model considers the management of a valued resource whose quality level changes over time. In the case of accumulating waste, for example, quality might be measured by the volume of waste: the smaller the amount, the higher the level of environmental quality. We denote quality at time t by a real number x_t , with larger values of x_t representing more desirable states. For mathematical convenience, we shall be working mostly with negative values for x . (Note that we use “quality” to denote the state of the resource. How the quality of a resource is *valued* will be captured in the utility function.) To keep things simple, we assume for the time being that there is a “manager” of the resource, who implements abatement and restoration policies in order to maximize the expected net present value of social welfare.³

We model the deterioration of the resource, absent the efforts of the resource manager, as a random variable with drift. The stochastic formulation reflects the reality that stock levels are not entirely in the control of the resource manager. Consider the case where environmental quality deteriorates as a stock pollutant accumulates. Although polluting firms are typically able to control their average emissions rates, the actual flow of pollution is likely to vary over time due to random shocks to firms’ inputs and outputs. For example, emissions of carbon dioxide from electric power plants vary over time with electricity generation, the type and quality of the fuels, the performance of the boilers, and so on – in ways that are neither entirely predictable, nor under the complete control of the plant operators. Even if emissions are controllable, their effects on the environment may not be: pollution concentrations depend on physical processes and climatic variables, not merely anthropogenic emissions.⁴ A similar role for stochastic variation exists in other settings. The deterioration of highways and bridges, for example, depends on weather conditions as well as on traffic levels and freight loads, whose realized values (at least from the point of view of the manager) are appropriately modeled as random fluctuations around a mean.

To capture such randomness in a simple way, we assume that cumulative deterioration up to time t , denoted by z_t , follows a Brownian motion with drift rate $\mu > 0$, variance rate σ^2 , and $z_0 = 0$. Hence, deterioration evolves according to $z_t = \mu t - \sigma w_t$, where w_t follows a standard Brownian motion. Unless the manager curbs the rate of deterioration or restores the resource, therefore,

³The manager could be the administrator of a regulatory agency that issues rules or provides rewards to influence the behavior of private-sector firms. We focus here on the behavior that a central planner would prescribe, and defer issues of instrument design to the discussion in Section 4 below.

⁴For example, the formation of ground-level ozone (a local air pollutant) depends on a complex interaction between nitrogen oxide emissions and biogenic volatile organic compounds, and is highly sensitive to temperature and sunlight.

quality at time t will be

$$x_t = -\mu t + \sigma w_t. \tag{1}$$

Intuitively, μ can be thought of as the “average” rate of deterioration of the resource: for example, average pollution emissions minus natural decay. Throughout the analysis, we will refer to the drift rate μ as the “flow rate,” and will use the terms “flow” and “deterioration” interchangeably. The relative values of the flow rate (μ) and the variance rate (σ^2) will vary with the setting. In contexts where natural recovery is negligible – e.g., a bridge – σ^2 will be small relative to μ .

2.1 Utility and cost functions

We assume that society’s benefit from the resource at any point in time depends only on the level of quality. Thus, at time t society derives a flow of utility $u(x_t)$ from the availability of the resource.⁵ The social rate of time preference is denoted by $\alpha > 0$. 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 , where E_x denotes the expectation conditional on an initial state x .*

Note that utility takes negative values; the utility function can be thought of as the negative of a convex loss function.

We define abatement as a reduction in the rate of deterioration: abating at rate a slows the expected deterioration rate from μ to $\mu - a$. Crucially, its costs are increasing on the margin.

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 $\epsilon > 0$.*

We assume that a finite maximum feasible rate of abatement exists, denoted \bar{a} .⁶ This ceiling may be higher than the mean flow rate μ . Hence our model allows (but does not impose) the

⁵We ignore issues such as population growth or changes in income, which could make the utility function time-dependent. For example, one might scale the utility function to the size of the population. If abatement costs remained constant while population grew, the optimal level of abatement at a given level of quality would increase over time. On the other hand, abatement costs and the drift rate μ might be greater for a larger population.

⁶The assumption of a ceiling on abatement provides a measure of generality. In some cases of interest, the manager may have limited abilities to stem or particularly to reverse the flow of deterioration. This assumption is completely innocuous, however: the ceiling can always be set high enough that the probability it binds is vanishingly small. Moreover, an optimal abatement policy can still be shown to exist even if we allow abatement to be unbounded.

possibility that abatement may more than fully offset deterioration. In such a case, “abatement” results in a positive rate of change in quality – but with increasing marginal costs.

Restoration corresponds to an improvement in quality that affects the stock directly, rather than by slowing deterioration. Two simplifying assumptions ease exposition; neither is crucial to our results. First, the manager can restore the resource from any state x_t to a certain high level, which we normalize as $x = 0$. (We relax this assumption of a fixed destination in Section 3.4.2, below.) Second, there is a positive fixed cost of restoration to $x = 0$, with zero marginal cost to starting at a lower point. This is an extreme form of nonconvexity.

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

Thus the cost of restoration is “destination-driven” in the sense of Phillips and Zeckhauser (1998): it depends on the ultimate level of quality, rather than the initial level (or the size of the quality gain). Though this assumption simplifies the model, our results hold for cost functions exhibiting less extreme economies of scale, as we discuss below in Section 3.4.2. For simplicity, we also assume that the restoration cost is time-independent.⁷

2.2 Abatement and restoration policies

An abatement policy $a(x)$ specifies the abatement level as a function of the state x . Under a restoration policy R , restoration occurs whenever the state x_t lies in the set R .⁸ Given a combined abatement–restoration policy (a, R) , the state of the resource evolves according to $x_t = \int_{s=0}^t (a(x_s) - \mu) ds + \sigma w_t - \sum_{\{i|T_i < t\}} x_{T_i}$, where T_i is the time at which the i^{th} restoration occurs. Hence starting from an initial state x , the infinite-horizon expected discounted utility is

$$E_x^{a,R} \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt - \sum_{i=1}^{\infty} e^{-\alpha T_i} C \right]. \quad (2)$$

⁷Two sources of time dependence would seem to be of potential interest. First, technological change could drive down the cost of restoration over time, raising the restoration trigger. Of course, technological advance in abatement would have a countervailing effect. A more interesting extension might have the restoration cost depend on the number of previous restorations, as we discuss in footnote 29.

⁸More formally, an *abatement policy* is a mapping $a : \mathfrak{X} \mapsto [0, \bar{a}]$, while a *restoration policy* is characterized by a measurable closed subset R of \mathfrak{X} . We shall restrict our attention to optimal stationary policies, but this does not affect the practical implications of our analysis. Suppose the manager’s problem is to choose an optimal stochastic process $\{a_t\}$ measurable with respect to the filtration generated by $\{w_t\}$. Such an optimal process can be produced by letting $a_t = a(x_t)$, justifying our approach.

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

3 Optimal abatement and restoration policies

In this section we characterize optimal restoration and abatement policies. We then discuss how the optimal policies vary with the magnitude and variability of flows, the costs of restoration and abatement, and the discount rate. Finally, we briefly consider three extensions of the basic model.

3.1 Characteristics of the optimal policies

We use stochastic dynamic programming to characterize the optimal restoration and abatement policies. Let J be the optimal value function:

$$J(x) = \sup_{a,R} E_x^{a,R} \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt - \sum_{i=1}^{\infty} e^{-\alpha T_i} C \right], \quad (3)$$

where the supremum is taken over pairs of abatement and restoration 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 describes the optimal abatement and restoration policies, and the resulting path of quality. It identifies two key states: \underline{x} , the restoration trigger; and x^\dagger , an inflection point in the value function that coincides with maximum abatement. (All proofs are given in the Appendix.)

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 results hold:*

Qualities of the value function : (i) $J < 0$ and $J(x)$ is finite for every x . (ii) $J(x) = J(0) - C$ for all $x \leq \underline{x}$. (iii) J is continuously differentiable for every x , and is twice continuously differentiable on (\underline{x}, ∞) . (iv) For all $x > \underline{x}$, J satisfies

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

Shape of the value function : (v) $J'(x) > 0$ for $x \in (\underline{x}, \infty)$. Moreover, (vi) $J''(x) > 0$ for $x \in (\underline{x}, x^\dagger)$; (vii) $J''(x^\dagger) = 0$; and (viii) $J''(x) < 0$ for $x \in (x^\dagger, \infty)$.

Optimal policy : (ix) 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). (x) a^* is increasing on $\{x \in (\underline{x}, x^\dagger) | a(x) \neq \bar{a}\}$ and decreasing on $\{x \in (x^\dagger, \infty) | a(x) \neq \bar{a}\}$. (xi) Letting $R^* = (-\infty, \underline{x}]$, the pair (a^*, R^*) is an optimal policy.

Under the optimal restoration policy, the manager restores the resource whenever quality falls to \underline{x} . This closely 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 quality, and a restoration corresponds to a replenishment of inventory.

The optimal abatement policy can be understood heuristically as equating marginal benefit and marginal cost at each level of quality. From Theorem 1, the abatement rate must attain the supremum of a function $f_x(a) = aJ'(x) - c(a)$, i.e., the components of equation (4) that are a function of a . The first term, $aJ'(x)$, represents the rate at which the value function increases. This corresponds roughly to the expected *benefit* from abating at rate a , taking into account present and future utility.⁹ The second term, $c(a)$, represents the cost of abatement a . Hence the optimal policy at each state sets the abatement rate to maximize the resulting “expected net benefit.”

Figure 1 illustrates the resulting abatement path.¹⁰ Because utility is concave, the marginal benefit from abatement at first increases as x diminishes. Thus starting from an initial high level (i.e., immediately following a restoration), the optimal abatement rate rises as quality worsens. Marginal benefit, and hence abatement, reach their peak at a state $x^\dagger > \underline{x}$.¹¹ Beyond this point, abatement *decreases* as quality continues to worsen. As the trigger point \underline{x} nears, the marginal benefit of abatement diminishes, since the quality of the resource will soon be restored.¹²

⁹Heuristically, for a given marginal change in the state dx , the resulting change in the value function would be $J'(x)dx$. Abatement a , carried out over an infinitesimal time period of duration dt , yields a marginal improvement in the state due to abatement $dx = adt$. We can think of $(adt)J'(x)$ as the resulting change in the value function (over an infinitesimal period of time). Dividing through by dt yields the rate of change in the value function, $aJ'(x)$. (Note that this heuristic explanation, like subsequent ones, puts intuition ahead of rigor, and thus is less technically precise than the formal results it seeks to explain.)

¹⁰The functional forms and parameter values used for all figures are provided in Appendix B.

¹¹The key state x^\dagger marks an inflection point in the value function: J is concave above x^\dagger and convex below. The convexity of the value function below x^\dagger – despite the concavity of the underlying utility function – is a consequence of the optimal restoration policy. J 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. The upper bound of this region is the inflection point x^\dagger .

¹²For destination-driven costs, the abatement rate falls to zero at the restoration trigger \underline{x} . At that point, the marginal benefits of further abatement are zero, because the state will be restored immediately. By the smooth-

If the abatement rate rises high enough, it will equal or exceed the average flow rate μ . In this case, let x^* denote the highest state at which $a(x^*) = \mu$. At states just below x^* , abatement is greater than average deterioration: hence the quality level will increase toward x^* in expectation. Above x^* , abatement slackens, and quality tends back toward x^* . These dynamics are illustrated by the shaded arrows in Figure 1. Note that $c'(\mu) = J'(x^*)$: thus at x^* , the marginal cost of fully abating expected pollution just equals the marginal benefit from doing so. This equimarginal condition suggests a useful analogy between x^* and the steady-state equilibrium in deterministic models of resource stocks.

We refer to x^* as the *mean drift equilibrium*. By this we mean that the optimal abatement rate at x^* equals the mean drift rate μ ; hence the expected rate of change in the quality of the resource is zero. Under the optimal policy, moreover, quality tends back toward x^* (in expectation) from states in the neighborhood of x^* . Of course, since flows are stochastic in our model, quality will not remain at x^* . Indeed, if the state deviates downward sufficiently (an event which will occur with certainty after a long enough period, given the stochastic process), restoration will be undertaken. This influences the optimal abatement path, as we discuss in the next section.

3.2 The interdependence of abatement and restoration

While the optimal restoration and abatement policies share features with familiar models, neither strategy takes the form it would in the absence of the other. The effect of abatement on restoration is intriguing. Intuition might suggest that restoration and abatement should substitute for one another: thus one might expect the optimal restoration trigger to fall when abatement is available. However, abatement can also raise the trigger. By allowing the manager to maintain a higher level of quality than she otherwise could, abatement may make high-quality states more attractive relative to low-quality ones. If so, the restoration trigger will increase, because the state is restored as soon as the gain in the value function equals the restoration cost. Simulations demonstrate that this latter effect dominates when flows are sufficiently great.

The effect of restoration on abatement, on the other hand, is unambiguous: less abatement is optimal when restoration is feasible. The feasibility of restoration must raise $J(x)$ everywhere, since its absence represents a constraint on the resource manager. But the value function increases

pastoring condition (Krylov, 1980), the marginal benefits from abatement must decline smoothly to zero as the state approaches \underline{x} . Marginal cost, and hence abatement, must follow suit. In section 3.4.2 below, we consider a variable component of restoration cost $\gamma(x)$, with $\gamma'(x) < 0$. In that case, abatement would decline smoothly to $\underline{a} > 0$ satisfying $c'(\underline{a}) = -\gamma'(\underline{x})$.

more at low levels of quality, where restoration is imminent, than at high levels, where restoration is more distant. Since $J'(x)$ is smaller when restoration is possible, the marginal benefit of abatement is also lower, hence less abatement is optimal. Theorem 2 states this result formally.

Theorem 2 *Let Assumptions 1, 2, and 3 hold. Let J_{abate} and a_{abate} denote the optimal value function and abatement policy in the absence of restoration. Then (i) $J' < J'_{abate}$; and (ii) for each state $x \in (\underline{x}, \infty)$, where \underline{x} is the restoration trigger, either $a(x) < a_{abate}(x)$ or $a(x) = a_{abate}(x) = \bar{a}$.*

Figure 2 portrays optimal abatement policies with and without the possibility of restoration. In the top panel, a mean drift equilibrium exists in both cases, but it occurs at a lower level of quality when restoration is available. Thus, restoration alters the optimal abatement policy even when it is exceedingly rare.¹³ In the bottom panel of the figure, restoration has a more drastic effect. No mean drift equilibrium exists: at *all* values of x above \underline{x} , abatement merely slows – but never halts – the net flow of damages. Rather than maintaining quality at a certain level, the optimal policy lets damages accumulate steadily until the trigger level is reached and the resource is restored. Which of the two cases portrayed in Figure 2 prevails depends on how rapidly the resource deteriorates, as we discuss in the next section.

3.3 The optimal balance of strategies

While the optimal policy always employs both restoration and abatement, its reliance on one strategy versus another depends on the rate and variability of flows and on economic variables such as costs and the discount rate. In this section, we consider how these factors affect the optimal mix of the two methods for improving quality. We rely on simulations for most of our results, since important relationships in the model are often complex, hence resistant to straight analytic demonstrations.

3.3.1 Mean flows and variability

First, consider the effects of the mean flow rate μ . Figure 3 illustrates optimal abatement policies (both in absolute terms and as a fraction of flows) for three flow rates. When the mean flow rate

¹³We thank an anonymous referee for pointing out a parallel to the optimal extraction of an exhaustible resource with an uncertain backstop technology. Even a small probability that such a technology will be developed in the future raises the optimal extraction rate (Dasgupta and Heal 1974; Dasgupta and Stiglitz 1981). In our model, a small probability of a future restoration lowers the optimal abatement rate. The decrease in abatement, speeding up deterioration, is akin to an increase in extraction that more rapidly depletes the stock.

is high, the cost of offsetting it with abatement is high as well. At the same time, restoration will be more frequent, on average, so that damages will persist for a shorter period before the resource is restored. Hence at higher flow rates, restoration becomes more attractive relative to abatement, and less abatement is done.¹⁴ The frequency distributions of states corresponding to these three flow rates are plotted in Figure 4. When flows are low, restoration is rare. States around the mean drift equilibrium x^* are much more common than other states, yielding a peak in the frequency distribution. At higher flow rates, the frequency distribution is flatter.

The fraction of total quality improvement achieved by abatement provides a natural way to describe the balance between the two strategies. Figure 5 plots this measure of abatement's importance against the flow rate. The dashed line on the figure marks the flow rate above which a mean drift equilibrium ceases to exist. (In Figure 5, this critical value is just above 1.2.) When flows are low, the cost of fully offsetting expected deterioration is also small. In this case, a mean drift equilibrium is reached, and restoration occurs with very small probability. For low flow rates, then, abatement is the principal improvement strategy. When flows increase beyond this cutoff, the mean drift equilibrium vanishes, and restoration becomes the principal strategy.

Greater variability has a similar effect as greater flows. Figure 6 depicts the effects of the variance rate σ^2 on the optimal abatement policy. Note that abatement reaches a higher peak when variability is lower. Given convex abatement costs, expected abatement cost increases when variability rises; hence less abatement is done. For sufficiently high variance rates, peak abatement never exceeds the mean flow rate, hence no mean drift equilibrium exists. Plotting the fraction of quality improvement due to abatement against the variance rate σ^2 would reveal a pattern similar to Figure 5. Since greater variability of flows depresses the optimal abatement rate, the share of improvement achieved by abatement falls as the variability rises.

For the range of parameters depicted in Figure 6, greater variability drives down the restoration trigger. This observation suggests a useful analogy with the theory of real options.¹⁵ Recall from Theorem 1 that the value function is convex just above the restoration trigger, and flat below it. Since the rate of deterioration is stochastic, and restoration is an irreversible investment, there is an option value to waiting before restoring. Just above the trigger, a favorable shock raises quality and expected utility. The "downside risk," however, is limited, since quality does not affect the

¹⁴This relationship can be neatly summarized in a poem: "Abate, don't wait, when flows are low / If flows are more you must restore."

¹⁵See Dixit and Pindyck (1994) for a thorough discussion of option value in the context of dynamic stochastic models of investment under uncertainty.

cost of restoration. The resulting option value represents a reward to waiting, and thus lowers the restoration trigger. As with options, the reward is greater when flows are more variable.¹⁶

3.3.2 Costs and discounting

Economic variables have effects that accord with intuition. When marginal abatement cost is higher, the optimal abatement rate is lower at every state, and the fraction of quality improvement achieved by abatement falls. When restoration is more expensive, more abatement is done. These effects are summarized in Figure 7, which demonstrates how the costs of abatement and restoration affect the amount of improvement done by each. For the purposes of illustration, and in keeping with our other numerical simulations, we depict the case of an exponential utility function; for that case, optimal policy is independent of the scale of abatement and restoration costs, so that their relative importance in promoting resource quality depends only on their relative cost. For three flow rates, the figure identifies two salient ratios of restoration cost C to the abatement cost parameter γ (which scales the slope of the marginal abatement cost function). When C/γ lies below the lower cutoff, abatement optimally accounts for less than five percent of the improvement in quality. Equivalently, the average abatement rate is below one-twentieth of the average flow rate in this region. When C/γ is above the upper cutoff, more than 95% of quality improvement is due to abatement.

The effect of higher discount rates, depicted in Figure 8, is striking. When quality is high, more patience leads to more abatement. This is intuitive, since the costs of abatement are incurred immediately, while the benefits stretch into the future. At low quality, however, the relationship inverts. A more patient manager will restore the resource at a higher quality, and hence abates less. That is because a lower discount rate also makes restoration more attractive, since the future benefits from a restored resource are valued more highly. In conventional resource management models, patience always prompts immediate preservation, e.g., harvests from a depleted fishery are reduced in the near term to build up the steady-state stock. In the SFQ model, by contrast, over a range of states a more patient manager will be more tolerant of short-run degradation – because such tolerance hastens restoration, and raises environmental quality in the long run.

¹⁶Alas, this relationship breaks down when both mean flow and variability are sufficiently low. We conjecture the following: When flows are sufficiently high that a mean drift equilibrium does not exist, the restoration trigger decreases monotonically with the mean flow rate and with the variability of flows.

3.3.3 The value of variance

Finally, consider how the variance rate σ^2 affects the value of the resource. Because utility is concave in quality, and abatement costs are convex, initial intuition might be that variance is always undesirable. Absent restoration, this would be true, since the value function would be everywhere concave, and Jensen’s Inequality would imply that variability lowers value. But when restoration is possible, J is convex immediately above the restoration trigger. In that region, variance increases expected net benefits.

Hence when restoration is feasible, a sufficiently patient manager may find variance to be desirable. The intuition is as follows. Consider accumulated deterioration z_t (as distinct from current quality x_t , which incorporates the effects of restorations). For example, z_t might represent total emissions (gross of abatement) of a stock pollutant up to time t . Since flows are stochastic, accumulated deterioration rises and falls; but it must always be less than or equal to the *maximum* cumulative deterioration that has been realized. The difference between the historical “high-water mark” and the current total must increase with the variability of flows. Given accumulated deterioration up to time t , therefore, the number of past restorations is greater when flows are more variable.¹⁷ Hence the quality of the state at time t will be higher, on average. If the discount rate is sufficiently low, this exchange of more numerous cleanups for higher quality raises the present value of expected utility.

We can establish a strong result for the case when only restoration is possible. In such a scenario, when the discount rate α is sufficiently small, variance raises the value function everywhere.

Theorem 3 *Let Assumptions 1 and 3 hold. Let $J_{restore}(\cdot, \sigma^2, \alpha)$ be the optimal value function given variance rate σ^2 and discount rate α , when only restoration is feasible. Then, for any x , there exists a scalar $\bar{\alpha} > 0$ such that for any $\alpha \in (0, \bar{\alpha})$, $J_{restore}(x, \sigma^2, \alpha)$ is increasing in σ^2 .*

Although Theorem 3 contemplates a polar case, it identifies a desirable aspect of variance that persists in the full SFQ setting. Because variability also has unwelcome effects when abatement is available, the value of variance is ambiguous in general. We leave a full analysis to future research.

¹⁷Initial intuition can be tricky here. Greater variance does not increase the expected frequency of restorations, which depends rather on the mean flow. However, at any particular point in time, the number of *past* restorations, conditional on cumulative damages z_t , increases with the variability of flows σ^2 .

3.4 Three extensions

3.4.1 Initial stock of quality

For convenience, in describing the optimal abatement policy, we have assumed that the starting point at time 0 coincides with the destination of restoration, i.e., that $x_0 = 0$. Of course, in real-world applications one might be interested in another starting point. Initial conditions do not affect the optimal abatement or restoration policies, which are functions of the state. Suppose, for example, that the resource manager only becomes aware of the deterioration in quality (or only becomes empowered to act) once the stock has fallen below the restoration trigger \underline{x} , as has happened with toxic waste dumps and endangered species in the U.S. In that case, immediate restoration would be optimal.¹⁸ On the other hand, if the initial state is relatively pristine, so that $x_0 > 0$, the optimal abatement rate will increase smoothly from a low starting point as quality falls toward 0, whence it proceeds as already described.

3.4.2 Greater costs for greater restoration

We have assumed that restoration costs are “destination-driven,” in that they depend on ultimate rather than initial quality. However, the results of the model hold for cost functions that exhibit less extreme economies of scale. For example, suppose that the cost of restoring the resource to $x = 0$ starting from quality level x has a fixed component F , as before, but also has a variable component $\gamma(x)$. We suppose $\gamma(x)$ to be a decreasing function of x ; i.e., the restoration cost increases with the amount of restoration. Total cost is given by $C(x) = F + \gamma(x)$. Unless $J(0) - J(x) < C(x)$ for all x , restoration will be optimal for at least one state x . Suppose such a state exists, and let \underline{x} denote 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 for restoration.¹⁹

Note that in this case, there is no longer any assurance that restoration will play any part of an optimal scheme. When restoration costs are destination-driven, restoration is guaranteed to

¹⁸In a previous version of this paper (Keohane, Van Roy, and Zeckhauser 2005), we studied the optimal management policy for the endangered California condor. Our simulations suggested that restoration would optimally occur at a population of about 160 birds. In fact, the number of condors had already fallen to a few dozen by 1967, when the species was first designated as “endangered”. In the mid-1980s, after the population had dwindled further, the U. S. Fish and Wildlife Service captured the remaining birds and embarked on a major restoration effort through captive breeding.

¹⁹If $\gamma(x)$ is convex, the restoration policy R may no longer be a convex set. Nonetheless, the evolution of the system will be similar, since restoration will be triggered each time the state hits \underline{x} . In this case too, the extent of scale economies will clearly depend on the relative size of the fixed cost F , and will determine how far the state falls before restoration is undertaken.

be attractive once quality deteriorates to a sufficiently low level. In contrast, when restoration is costly at the margin, abatement may always dominate restoration even at low states. In particular, if the variable cost of restoration is sufficiently high, then $J(0) - J(x) < C(x)$ for all x . In that case, restoration will never occur: instead, quality will be maintained only by abatement, and the state will fluctuate around the mean drift equilibrium x^* . In a sense, this represents a limiting case of our analysis. The distinguishing feature of restoration in our model is the presence of scale economies.

We can extend this discussion further to allow the destination of restoration to be endogenous. Consider a model in which restoration can result in any quality level $x \in [x_t, \infty)$, and the cost incurred is a function C of the quality difference $(x - x_t)$. Let restoration cost increase with the amount of restoration done, but at a decreasing rate (i.e., exhibit economies of scale); thus $C > 0$, $C' > \epsilon$ for some positive scalar ϵ , and $C'' < 0$. In this case, as in the earlier ones, it can be shown that there exists a state \underline{x} such that restoration occurs if and only if $x_t \leq \underline{x}$. Unlike the previously discussed cases, the optimal destination state $\bar{x}(x)$ generally will depend on the state x at the time of restoration. Because $C'' < 0$, $\bar{x}(x)$ increases as x decreases. Together with the fact that x_t is continuous, this implies that in steady state, restorations only occur at \underline{x} and quality is restored to $\bar{x}(\underline{x})$ each time this happens. In particular, Theorem 1 holds so long as Assertion (ii) is replaced by the following: For each $x \leq \underline{x}$, there is a state $\bar{x}(x) > x$ such that $J(x) = J(\bar{x}(x)) - C(\bar{x}(x) - x)$; furthermore, $J(x)$ is increasing and $\bar{x}(x)$ is decreasing on $(-\infty, \underline{x}]$.

3.4.3 Delayed restoration

In many real-world applications of the SFQ model, restoration is unlikely to be instantaneous. For example, proposed methods to remove greenhouse gases from the atmosphere (e.g., seeding oceans with iron filings to promote plankton growth) would require long lead times. Consider generalizing the model of Section 2 to incorporate a delay of length D in restoration. During the interval $[T_i + D, T_{i+1})$ between the completion of the i^{th} restoration and the commencement of the next, the state evolves as before. The optimal value function is now defined by

$$\begin{aligned}
 J(x) = & \sup_{a,R} E_x^{a,R} \left[\int_{t=0}^{T_1} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right. \\
 & \left. + \sum_{i=1}^{\infty} \int_{t=T_i+D}^{T_{i+1}} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt - \sum_{i=1}^{\infty} e^{-\alpha T_i} C \right]. \tag{5}
 \end{aligned}$$

(Compare (5) to (3).) Note that states observed during a restoration project do not enter into equation (5). Instead, the restoration cost C incorporates all costs incurred and utility realized in the course of a restoration project. While the length of the delay is deterministic in this model, the restoration cost could be random, in which case C would be the expected cost. The results in Theorems 1 and 2 continue to hold in such a model, with only a slight modification to Theorem 1: the boundary condition stated in Assertion (ii) becomes $J(x) = e^{-\alpha D} J(0) - C$ for all $x \leq \underline{x}$.

4 Policy implementation when flows are decentralized

Thus far, we have assumed implicitly that the resource manager is able to implement the optimal abatement and restoration policy directly. In many applications of interest, however, no single “manager” exists. For example, carbon dioxide is emitted by factories, power plants, and motor vehicles; municipal solid waste is generated by firms and individuals; passing cars and trucks punch potholes in highways. To achieve optimality, a central authority (usually the government) must regulate the activities of these private parties.

In this section, we suppose that deterioration (and thus abatement) depends on a large number of self-interested agents who are subject to regulation by a central authority; and that authority also carries out any restoration. Two issues arise. First, how can the optimal abatement path be achieved? Second, how should the central authority support the costs of restoration? As a motivating example, we consider a stock pollutant emitted by a multi-firm industry. Let μ_j denote firm j 's average rate of emissions in the absence of any regulation, and denote firm j 's abatement rate at time t by a_{jt} . Thus the expected deterioration in quality (i.e., expected emissions net of abatement) is $\mu - a_t$, where $\mu \equiv \sum_j \mu_j$ and $a_t \equiv \sum_j a_{jt}$. Aggregate abatement cost is then $c(a_t) = \min_{\{a_{jt} | \sum_j a_{jt} = a_t\}} \sum_{j=1}^N c_j(a_{jt})$. For simplicity, we shall assume that firms are “small” relative to the industry, so that each firm ignores the effects of its own emissions and abatement on the state.

4.1 Aligning incentives for abatement

As in other settings, the central authority can induce efficient abatement by levying an appropriate tax on pollution, or more generally on deterioration. Consider a tax $\tau(x)$ set equal to the derivative of the value function, $J'(x)$, which represents the present discounted value of the marginal benefit of abatement. Faced with such a tax, each polluter that chooses positive abatement will equate

its marginal abatement cost with the tax.²⁰ Hence $c'_j(a_{jt}) = \tau(x) = J'(x)$, ensuring that at every state $x > \underline{x}$ the optimality condition $c'(a) = J'(x)$ is met in equation (4). Note that while the tax varies over time to equal marginal damages, it is linear in emissions at every state x . Thus at any given point in time, each polluter faces a constant tax per ton of emissions, and all polluters face a common tax rate; but this tax rate changes over time with the quality of the resource.

Several characteristics of the optimal tax are notable. First, one can readily show that the expected tax revenue to the center is the same whether it collects the flow tax $\tau(x)$ on the actual change in the state $(-dx_t)$, or on the expected deterioration over a given time interval, $(\mu - a_t) dt$.²¹ Abatement incentives are also identical in the two cases. If firms are risk averse, the center may prefer to charge them only for expected deterioration, so that it (rather than the firms) bears the risk from random fluctuations. Second, to align incentives properly, the regulator must both levy a tax on net emissions and pay out a subsidy on net abatement. Thus the expected flow of revenue to the regulator, $\tau(x)(\mu - a_t)$, will be positive or negative depending on $\mu \gtrless a_t$.²²

Third, the optimal tax will follow the same qualitative path as abatement. Indeed, a plot of the tax rate against quality would be identical to panel (a) of Figure 3, only scaled up by a constant.²³ When environmental quality is high, the present value of marginal damages is low, for two reasons: current marginal damage is low because the stock of pollution is small; and the contribution of emissions to future high-damage states is far removed, hence heavily discounted. As quality deteriorates, the tax rises along with the present value of marginal damages. Once quality has deteriorated sufficiently, however, the optimal tax falls: as restoration nears, additional pollution will affect future utility over a shorter duration (at least in expectation), and thus impose smaller marginal damage.

The eventual decline in the optimal tax rate contrasts with the steadily rising taxes derived in most theoretical models of stock pollution (Plourde 1972; Falk and Mendelsohn 1993; Farzin 1996),

²⁰The perfect competition assumption is crucial here. If firms are sufficiently large that their abatement affects the state, they will have incentives to over- or under-abate, relative to the optimal path, in order to lower their future taxes. For example, over the range of states when the tax rate is decreasing in x (i.e., when J is concave), strategic firms would abate more than the efficient rate, in order to slow the rate at which quality deteriorated and their taxes increased. Below the inflection point x^\dagger , on the other hand, such firms would abate less than the efficient rate in order to hasten the arrival of lower quality states with lower taxes.

²¹This equivalence follows from the fact that $E \left[\int_{t=T_{i-1}}^{T_i} \tau(x_t) \sigma dw_t \right] = 0$ by the optimal sampling theorem: note that $\int_{s=T_{i-1}}^t \tau(x_s) \sigma dw_s$, conditioned on T_{i-1} , is a martingale, and T_i is a stopping time.

²²Positive net abatement, $a > \mu$, would be feasible in cases such as global climate change, where afforestation allows negative net flows of carbon. Of course, in many cases – e.g., soot from factories – it is natural to assume that $\bar{a} \leq \mu$, so that abatement cannot exceed unregulated emissions. If so, the expected tax revenue will always be positive.

²³The scaling factor is the slope of the marginal cost function, equal to 80 in our simulations.

as well as integrated assessment models of climate change (Nordhaus and Boyer 2000).²⁴ Analyses of an optimal carbon tax, on the other hand, have found a “hump-shaped” time profile similar to that of the SFQ tax (Ulph and Ulph 1994; Hoel and Kverndokk 1996; Farzin and Tahvonen 1996). In both the SFQ model and these optimal carbon tax models, the eventual fall in the optimal tax happens because the pollution stock ultimately gets eliminated. As the stock falls to zero, the marginal damages follow suit. The reason the stock disappears, however, is fundamentally different in the two types of models. In the carbon-tax literature, the stock pollutant (i.e., carbon dioxide) is modeled as the byproduct of consuming an exhaustible natural resource (i.e., fossil fuels). The pollution stock decays to zero once the scarce resource is exhausted, eliminating the source of emissions. In the SFQ model, the pollution stock vanishes because it is cleaned up.²⁵

Fourth, the optimal tax in our model bears a surprising relationship to the mean flow rate, μ . Because the optimal tax rate is a monotonic function of abatement, it reaches a lower peak when flows are high on average (recall Figure 3). Hence when restoration is feasible, the optimal emissions tax may be lower when unregulated emissions are higher – that is, when there is more pollution (or potential pollution) around.²⁶ This result contrasts sharply with familiar models, in which the efficient tax increases with the baseline level of pollution. The availability of restoration makes the difference. At high levels of quality, restoration is remote, and the dynamic remains much the same as the no-restoration case: A higher flow increases marginal damages from current pollution (and hence the tax), because it speeds the arrival of states in which marginal utility is high. The situation reverses at low quality levels. Then the predominant effect is that higher flow hastens restoration, which terminates the effects of current pollution; hence marginal damage is lower. Optimal abatement, and the optimal tax rate, fall accordingly.

²⁴If restoration were not available, our model would produce an essentially identical result to those deterministic models. The optimal tax would rise over time as quality worsened, until the resource reached the mean drift equilibrium (x^*); after that point, the tax would (in expectation) be constant, varying only with stochastic changes in the resource quality around x^* . Note further that the initial *increase* in the tax rate over time depends on the assumption that the initial pollution stock is lower than its eventual level. In the deterministic framework, this is readily seen in the first-order conditions for maximization, although it is not always recognized in the literature. In our framework, this amounts to the assumption that $x_0 > x^*$. Recall our discussion on the role of the starting point x_0 (section 3.4.1).

²⁵We thank an anonymous referee for noting an analogy between our model and that of Farzin and Tahvonen (1996). They show that under certain conditions, the optimal extraction of an exhaustible resource may follow an “open-close-open” cycle – i.e., positive extraction, followed by a period of zero extraction, followed by a second phase of positive extraction until exhaustion. This echoes the optimal SFQ policy of abatement punctuated by periodic restoration.

²⁶Note that in our model – as in standard models of abatement – abatement cost is assumed to be only a function of the amount abated, and thus independent of unregulated emissions (here represented by μ).

4.2 Raising funds for restoration

Next, consider how the central authority should fund restoration. This would seem to pose a new issue: how should the cost of a restoration project be divided up among the agents responsible for the deterioration of the resource? In fact, no new cost-sharing rule is needed, because the revenue raised from the flow tax closely approximates the restoration cost. Indeed, as the variability of flows goes to zero, the total revenue collected from the flow tax, between any one restoration and the next, converges *precisely* to the cost of restoration.

To see this, recall that the cost of restoration, C , must equal the gain in the value function from restoration. Thus $C = J(0) - J(x) = -\int_{t=T_{i-1}}^{T_i} dJ(x_t)$, where T_i (as before) is the time of the i^{th} restoration. Applying Ito's Lemma, the integral becomes:

$$C = \left(-\int_{t=T_{i-1}}^{T_i} J'(x_t) dx_t \right) - \left(\frac{\sigma^2}{2} \int_{t=T_{i-1}}^{T_i} J''(x_t) dt \right). \quad (6)$$

The first term on the right-hand side of equation (6) equals the tax revenue from levying the flow tax $\tau(x) = J'(x)$ on the observed deterioration in resource quality. The sign of the second term depends on the sign of J'' over the realized path of the state $\{x_t\}$. Since the value function has both convex and concave regions, this second term may be positive or negative: thus the revenue from the flow tax may exceed or fall short of the cost of restoration. As $\sigma^2 \rightarrow 0$, however, this term vanishes, so that the restoration cost precisely equals the revenue from the flow tax. Even when the variance is low but positive, the difference between the restoration cost and the revenues from the flow tax alone is negligible. Thus when variability is relatively small, the flow tax suffices to raise the funds needed for restoration.²⁷

²⁷Simulation results indicate that for the parameter values considered throughout Section 3 of this paper (i.e., $\sigma^2 = 9$ and $\mu \in [1, 3]$) the discrepancy is on the order of one percent of the restoration cost or less. However, if σ^2 were dramatically larger relative to flow rate, say a variance of 100 for a flow rate of 1, the funding gap would rise to around 35% of the restoration cost.

A natural way to achieve exact budget balance is suggested by the second term in equation (6), under the additional assumption that firm-level emissions are independent. Consider a tax $\rho(x) = -\frac{J''(x)}{2}$ on the *variance* generated by each firm per unit time. Such a tax would yield total revenue of $-\frac{\sigma^2}{2} J''(x) dt$ over an interval dt – ensuring that the combined revenue from the variance and flow taxes would precisely equal the cost of restoration, for any realization of the stochastic process. Note that such a variance tax is not needed to align incentives; even if variance is taken to be controllable by polluting firms, it could be contracted upon, given our assumption of Brownian motion. (We thank an anonymous referee for pointing this out.)

5 Applications

Here we explore the implications of our theoretical model for managing real-world resource stocks. First, we briefly sketch the model’s application to a number of environmental problems, ranging from accumulating waste to zebra mussels.²⁸ We then discuss how the model would apply to public infrastructure – in particular, highway construction and maintenance.

5.1 Environmental quality

The accumulation of wastes at disposal sites or generating facilities is a canonical SFQ problem. Consider the management of municipal solid waste, for example. The environmental quality of a landfill site and the surrounding area diminishes as garbage accumulates. The flow of waste may be slowed through recycling, composting, or waste reduction. Eventually, the landfill is capped, the site is restored – perhaps becoming a park or recreation area – and quality returns to its initial high level.²⁹ In a typical scenario in the real world, waste diversion remains roughly constant over time, or changes only with changing preferences (e.g., a desire to increase levels of recycling) or prices (e.g., land becomes more expensive, or recycled materials become more valuable). Optimal waste management would vary the rate of abatement over time. Early in the life of a landfill, diversion should be relatively high, since the discounted expected damages from dumping garbage are high relative to the damages from waste arriving later. As the landfill nears capacity, diversion should drop, since the waste will impose damages only for the brief time until restoration.

Similar issues, on a different scale, are involved in the management of hazardous wastes. Consider the chemistry department at Harvard University.³⁰ The department’s laboratories accumulate a variety of toxic and reactive substances. Storing such substances on campus heightens health and fire hazards.³¹ Removing the wastes for permanent disposal – restoration in this context –

²⁸For further discussion of environmental applications, see Keohane, Van Roy, and Zeckhauser (2005).

²⁹With 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 could be extended to accommodate the multiple-site case by having restoration costs rise as we move to successively more expensive landfills. Such an approach would characterize restoration as using up a nonrenewable resource (finite landfill space); hence we can appeal to results from the theory of nonrenewable resources (Dasgupta and Heal, 1979; Hotelling, 1931). Abatement today would be influenced by the shadow price of future restorations.

³⁰We thank Henry Littleboy, Health and Safety Officer (for Harvard’s Faculty of Arts and Sciences Office of Environmental Health Services), who oversees 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.

³¹Of course, chemical waste storage and disposal are heavily regulated by the Environmental Protection Agency. For example, existing regulations prohibit the storage of waste longer than ninety days. At Harvard, the constraint does not bind: limited storage space makes more frequent collection necessary.

involves economies of scale, reflecting the fixed costs of labor and transportation. A 55-gallon drum of corrosive flammable liquids costs \$320 to ship; a single 5-gallon container, \$95. At least in principle, several methods exist to control the flow of lab waste generated: experiments could be altered or curtailed to conserve chemicals; technicians could exert greater effort to prevent spills; laboratories could manage their inventories more efficiently; or some fraction of the waste stream could be purified and reused rather than thrown away. For years, however, individual laboratories were not charged for disposal, and thus had little incentive to reduce their chemical use. Limited experience with a recently-imposed volume-based charge indicates that the use of chemical wastes is fairly inelastic, suggesting high costs of substantial abatement.

The sedimentation of reservoirs presents a very different application.³² The “stock” in this context is the capacity of the reservoir, which diminishes as sediment flows into the reservoir and accumulates. Retiring a dam and constructing a replacement constitute restoration, with destination-driven costs. The sediment flow can be abated by soil conservation, reforestation, and other measures in the catchment area; or sediment can be routed away from the reservoir. The common practice of letting sediment accumulate unchecked before retiring the dam – equivalent to a restoration-only policy – is almost surely suboptimal. At the other extreme, “sustainable management” that seeks to maintain an equilibrium by relying exclusively on controlling sediment flow, without periodic restoration (Palmieri, Shah, and Dinar 2001), is equally unlikely to be optimal.

The SFQ model applies naturally to the control of animal pests, such as zebra mussels (*Dreissena polymorpha*). These small freshwater mollusks were introduced to the U. S. accidentally, carried in bilge water of cargo ships. They 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 region is estimated to have cost as much as \$1 billion in the 1990s alone.³³

Methods that can prevent mussel settlement vary by location. In the pumphouses of power plants, mussels grow on walls, debris screens, valves, and pumps, obstructing the flow of water. Mechanical measures to remove them – physical scraping or “hydrolasing” with high-powered water hoses – involve high fixed costs: sending down a team of divers or even dewatering the pumphouse

³²The description of dam sedimentation and management draws on Palmieri, Shah, and Dinar (2001).

³³Personal communication, Charles O’Neill, Project Director, National Zebra Mussel Information Clearinghouse, New York Sea Grant.

(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, rendering mechanical removal infeasible, but chemical removal is possible. In such locations, both flow and stock controls are employed. Continuous low-level chlorination of circulating water is an abatement policy that inhibits the settlement of juvenile mussels. Periodic injections of high concentrations of chlorine or other biocides represent a restoration strategy used to kill off encrusted adult mussels.

5.2 Public infrastructure

The basic SFQ model applies to a range of public infrastructure projects. Consider highways, for example. Highway expenditures take a substantial chunk of public spending, most of it routed to building new roads rather than maintaining old ones. Across countries, maintenance appears to be underfunded relative to new investment, particularly in developing countries (Rioja 2003). In the U.S., total government expenditures on highways averaged \$122.5 billion annually over the period 1998-2002; half went to capital outlays (new investment), while a quarter was spent on maintenance (USDOD, 2004).

Recent studies use optimal growth models to consider how public funds should be allocated for investment and maintenance in infrastructure (Rioja 2003; Kalaitzidakis and Kalyvitis 2004). An SFQ approach incorporates restoration to represent the replacement and/or rehabilitation of existing infrastructure.³⁴ Suppose that a large number of infrastructure projects coexist, and that they share a common level of unchecked deterioration. Then the quality of infrastructure projects will be described by a frequency distribution of the type illustrated in Figure 4, while the optimal share of expenditure on maintenance versus restoration at any point in time will be described by Figure 5.

Note how the rate of unchecked deterioration helps to determine the optimal policy. The greater is economic activity, all else equal, the faster will infrastructure deteriorate, and the greater should be the expenditure on new investment (restoration) rather than maintenance. One implication

³⁴In models of infrastructure, deterioration is typically assumed to be proportional to the existing capital stock, whereas our formulation assumes that it is independent of the stock. However, this difference can be easily accommodated. In the infrastructure setting, our model can be considered to apply to “single projects” subject to maintenance and restoration. In a model of a large number of such projects, each of identical size, aggregate deterioration would indeed be proportional to the number of projects, i.e., the total capital stock.

is that the share of expenditures allocated to maintenance relative to new construction should be *higher* in less-developed countries (where economic activity is lower); the reverse pattern is observed.

An SFQ analysis also offers insight into the optimal financing of infrastructure projects. A tax on deterioration – for example, a tonnage tax on trucks to fund highway expenditures – could align private incentives with social welfare, as it simultaneously raises the revenue required to pay for restoration. Indeed, current gasoline taxes in the U. S., which help pay for highway construction, are (at least qualitatively) appropriate instruments. Moreover, implementation of a targeted deterioration tax would be relatively easy, given the existing structure of gasoline taxes, tonnage charges, and turnpike tolls.

6 Conclusion

In a wide range of settings, both stocks and flows can be controlled to improve the quality of a valued resource. If so, the SFQ model applies. Managing the resource entails abating the downward drift in quality and periodically restoring the stock. These strategies are interdependent. The optimal balance between them depends on the rate and variability of ongoing deterioration, the costs of the two strategies, and the discount rate. If flows are low enough or abatement is cheap enough, a “mean drift equilibrium” may be reached where abatement efforts just offset the expected deterioration of the resource. If so, abatement is the principal management tool, although the potential for restoration still lowers the optimal abatement rate. When deterioration is more rapid or more variable, when abatement is more expensive, or when restoration is less costly, the optimal policy relies more on restoration.

This model has broad relevance for the management of resource stocks in the real world. We have discussed a range of applications: the disposal of municipal solid waste and hazardous laboratory waste; the slowing of siltation in reservoirs; the control of pests such as zebra mussels; and the construction and maintenance of public infrastructure projects such as highways. Our analysis generalizes readily to the management of private stocks of physical and human capital. Thus, determining how steadily to service a piece of capital equipment, and how often to replace it, depends on the speed and variability of deterioration, not merely the costs of maintenance and replacement. Similarly, from the perspective of the firm, investment in human capital presents an SFQ problem. Workers age, tire, and burn out. In industries with rapid technological advance,

workers' skills quickly obsolesce. A firm can train its workers to maintain their productivity, but at some point it may replace its older workers, through layoffs or reassignments to less cutting-edge tasks.

Government will likely play a role in many SFQ problems, e.g., in controlling environmental quality. A time-varying tax can ensure efficient abatement in a decentralized setting. Here, too, abatement and restoration are linked: the tax that induces efficient abatement generates the revenues needed to pay for restoration. Moreover, government policy should incorporate the central lesson of the SFQ analysis: that stock and flow controls should be coordinated and implemented jointly when both are feasible. When restoration is an option, maintaining a resource stock at a constant level by abating flows, the usual government policy, will be more expensive than achieving the same present value of expected utility from quality, but allowing quality to vary over time. On the other hand, a policy relying solely on restoration will not only restore too frequently (since deterioration is unchecked), but may also allow quality to fall too far before each restoration (since the availability of abatement may raise the optimal trigger).

Alas, single-prong strategies are often employed inappropriately in the real world. For example, environmental policies towards municipal and hazardous waste tend to emphasize terminal cleanup and permanent storage (restoration) rather than slowing waste generation. In contrast, traditional regulation of water quality, as embodied in the Clean Water Act, entails curbing effluents rather than boosting environmental quality. Restoration efforts have been viewed as last-ditch measures borne of desperation: consider the Boston Harbor cleanup, or the imminent remediation of the Hudson River. While those specific restorations may have been long overdue, they do not signal the failure of earlier efforts at controlling pollution. Rather, our model suggests that the periodic cleanups they exemplify are part of optimal policies to manage resource quality.

Appendix A: Proofs of Theorems

Proof of Theorem 1

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

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 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 $J(x) = \sup_a E_x^a \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right]$. It is easy to see that 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 $x > \underline{x}$, establishing Assertion (ii).

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

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

for all $x > \underline{x}$. Hence, Assertions (iii) and (iv) are valid. It is easily verified by sample-path arguments that J is increasing on (\underline{x}, ∞) (Assertion (v)).

It follows from Assertions (ii) and (iii) 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 (vi) and (vii).

Now consider an optimal policy. Assertion (ii) implies that the restoration component of an optimal policy is given by $R^* = (-\infty, \underline{x}]$. Let a function f_x be defined for $x > \underline{x}$ by $f_x(a) = aJ'(x) - c(a)$. Note that $f_x'' = -c'' \leq -\epsilon$ for some ϵ . Hence, for any x , the supremum

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

is uniquely attained by some $a \in [0, \bar{a}]$. 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 (ix) and (xi). Moreover, 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 (x) follows.

We are left with the task of establishing Assertion (viii). 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 $[0, T]$, where T is given by

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

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 $t \in [0, T]$. It is easy to show that T is finite with probability one.

Define “sample costs” associated with the three processes:

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

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

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

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

$x_T^- \neq x^\dagger$. In this event, we have $x_T^- = x_T^+ = x_T > x^\dagger$, and the desired inequality follows directly from concavity of u and convexity of c .

The second case is when $x_T^- = x^\dagger$. 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 $\sup_{t \in [0, T]} |a^*(x_t^-) - a^*(x_t^+)| = O(\Delta)$ and $x_T^+ - x^\dagger = O(\Delta)$. We then have

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

where the second-to-last expression relies on the fact that $J''(x^\dagger) = 0$ and that $x_T^+ - x^\dagger = O(\Delta)$. 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,

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

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

Next, consider Theorem 2, which contrasts the full SFQ case with the case where only abatement is possible. Before proving the theorem, we establish some properties of the optimal abatement policy in the abate-only case.

Lemma *Let Assumptions 1 and 2 hold, and assume that restoration is not feasible. Let \tilde{a} denote the optimal abatement policy in this case. Then (i) there exists a state \hat{x} such that \tilde{a} is decreasing on (\hat{x}, ∞) and $\tilde{a}(x) = \bar{a}$ for all $x \leq \hat{x}$; (ii) $\lim_{x \rightarrow \infty} \tilde{a}(x) = 0$; and (iii) if $\bar{a} > \mu$, then there exists a state x_{abate}^* such that $\mu < \tilde{a}(x)$ for $x < x_{abate}^*$ and $\mu > \tilde{a}(x)$ for $x > x_{abate}^*$.*

Proof. Let \tilde{J} denote the optimal value function in the case when only abatement is possible:

$$\tilde{J}(x) = \sup_a E_x^a \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right], \quad (8)$$

where the supremum is taken over abatement policies. Let \tilde{a} be the corresponding optimal abatement policy.

Let \tilde{f}_x be defined by

$$\tilde{f}_x(a) = a\tilde{J}'(x) - c(a),$$

and let $\tilde{a}(x)$ be the value in $[0, \bar{a}]$ that uniquely attains the supremum of \tilde{f}_x . Along similar lines as in the proof of Theorem 1, one can show that $\tilde{J}'(x) > 0$, implying $\tilde{f}'_x(0) > 0$. Also recall that $\tilde{f}''_x = -c'' \leq -\epsilon$ for some ϵ . Consider the less constrained problem

$$\sup_{z \in [0, \infty)} \tilde{f}_x(z). \quad (9)$$

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

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

Recall that $\tilde{J} < 0$ and $\tilde{J}' > 0$, so that $\lim_{x \rightarrow \infty} \tilde{J}'(x) = 0$. Hence for any $z > 0$, there exists a state x such that $\tilde{f}'_x(z) < 0$, implying that $\lim_{x \rightarrow \infty} b(x) = 0$ and that Assertion (ii) 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 $\mu < \bar{a}$ by hypothesis, we have Assertion (iii). **q.e.d.**

Proof of Theorem 2

Proof. As a step toward establishing Assertion (i), we will show that $\tilde{J} < J$. It is easy to see that $\tilde{J}' \leq J'$. From Theorem 1, 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 principal of optimality (see, e.g.,

Krylov), we have

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

and

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

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

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

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

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

Now turn to Assertion (ii). Again, let \tilde{f}_x be defined by $\tilde{f}_x(a) = a\tilde{J}'(x) - c(a)$. 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 (ii). **q.e.d.**

Proof of Theorem 3

Proof. Without loss of generality, we will assume in this proof that $\sigma \geq 0$. Recall that damage evolves according to $z_t = \mu t - \sigma w_t$. Consider a fixed restoration threshold $\tilde{x} < 0$, which may or may not correspond to the optimal restoration strategy. We introduce some notation to facilitate our analysis. First, we denote the running maximum of damage by $m_t = \max_{\tau \in [0,t]} z_\tau$. The number of restorations carried out up to time t is $r_t = \lfloor -m_t / \tilde{x} \rfloor$. Given only knowledge of z_t , the tightest

lower bound on r_t is $\underline{r}_t = \lfloor -z_t/\tilde{x} \rfloor$. The state can be written as $x_t = -z_t - r_t\tilde{x}$. If we carried out \underline{r}_t rather than r_t restorations, the state would be $y_t = -z_t - \underline{r}_t\tilde{x}$.

Let $J_{\tilde{x}}(\cdot, \sigma, \alpha)$ be the value function corresponding to a restoration threshold \tilde{x} . Since x_t reaches \tilde{x} in finite expected time and the process regenerates every time it hits \tilde{x} , it is ergodic. It follows that

$$\begin{aligned} \lim_{\alpha \downarrow 0} \alpha J_{\tilde{x}}(x, \sigma, \alpha) &= \lim_{\alpha \downarrow 0} \alpha E_x \left[\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt + \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} E_x \left[\int_{t=0}^T u(x_t) dt + r_T C \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} E_x \left[\int_{t=0}^T u(x_t) dt \right] - C\mu\tilde{x}, \end{aligned}$$

where the final term follows from the fact that the expected interarrival time between visits to \tilde{x} is $-\mu\tilde{x}$.

We will now establish that $\lim_{\alpha \downarrow 0} \alpha J_{\tilde{x}}(x, \sigma, \alpha)$ is increasing in σ . Note that $(x_t, y_t, r_t - \underline{r}_t)$ together form an ergodic process. There is a joint stationary distribution over the variables x_t , y_t , and $r_t - \underline{r}_t$ such that if $(x_0, y_0, r_0 - \underline{r}_0)$ is sampled from this distribution, $(x_t, y_t, r_t - \underline{r}_t)$ is a stationary process. Let E_{∞} denote expectation with respect to the distribution of this stationary process. It is easy to see that, for any t , the marginal distribution (with respect to the stationary process) of y_t is uniform over $[\tilde{x}, 0]$. We therefore have

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} E_x \left[\int_{t=0}^T u(x_t) dt \right] &= \lim_{T \rightarrow \infty} \frac{1}{T} E_{\infty} \left[\int_{t=0}^T u(x_t) dt \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} E_{\infty} \left[E_{\infty} \left[\int_{t=0}^T u(x_t) dt \middle| y_t \right] \right] \\ &= \lim_{T \rightarrow \infty} \frac{1}{\tilde{x}T} \int_{y=0}^{\tilde{x}} E_{\infty} \left[\int_{t=0}^T u(y_t - (r_t - \underline{r}_t)\tilde{x}) dt \middle| y_t = y \right] dy. \end{aligned}$$

Note that, conditioned on z_0 and z_t , the process z_{τ} forms a Brownian bridge on $\tau \in [0, t]$. A sample path argument shows that for any $\gamma > \max(z_0, z_t)$, $\Pr\{m_t \geq \gamma | z_0, z_t\}$ is increasing in σ . It follows that for any $\gamma > \max(z_0, z_t)$, $\Pr\{m_t - z_t \geq \gamma | z_0, z_t\}$ is increasing in σ , and therefore, for any $\gamma \geq 1$, $\Pr\{r_t - \underline{r}_t \geq \gamma | z_0, z_t\}$ is increasing in σ . Since this holds for all z_0 and z_t , and y_t is a deterministic function of z_t , for any $\gamma \geq 1$ and any y_t , $\Pr\{r_t - \underline{r}_t \geq \gamma | y_t\}$ is also increasing in σ . Since $u' > 0$, it follows that

$$E_{\infty} \left[\int_{t=0}^T u(y_t - (r_t - \underline{r}_t)\tilde{x}) dt \middle| y_t \in dy \right] dy$$

is increasing in σ . Therefore,

$$\lim_{T \rightarrow \infty} \frac{1}{T} E_x \left[\int_{t=0}^T u(x_t) dt \right]$$

is increasing in σ . It follows that that $\lim_{\alpha \downarrow 0} \alpha J_{\tilde{x}}(x, \sigma, \alpha)$ is increasing in σ .

It is not hard to show that for any $x > \tilde{x}$ and any $\alpha > 0$, $J_{\tilde{x}}(x, \sigma, \alpha)$ is continuously differentiable in \tilde{x} and σ , and we will take this as given. Let $\underline{x}(\sigma, \alpha)$ denote the optimal threshold as a function of σ and α . It can be shown that $\underline{x}(\sigma, \alpha)$ is continuously differentiable in σ , and we take this as given as well. It follows that

$$\frac{\partial J_{\text{restore}}(x, \sigma, \alpha)}{\partial \sigma} \Big|_{\sigma=\bar{\sigma}} = \frac{\partial J_{\tilde{x}}(x, \bar{\sigma}, \alpha)}{\partial \tilde{x}} \Big|_{\tilde{x}=\underline{x}(\bar{\sigma}, \alpha)} \frac{\partial \underline{x}(\sigma, \alpha)}{\partial \sigma} \Big|_{\sigma=\bar{\sigma}} + \frac{\partial J_{\underline{x}(\bar{\sigma}, \alpha)}(x, \sigma, \alpha)}{\partial \sigma} \Big|_{\sigma=\bar{\sigma}}.$$

Since $\underline{x}(\sigma, \alpha)$ maximizes $J_{\tilde{x}}(x, \sigma, \alpha)$ over $\tilde{x} \in \mathfrak{R}$, we have

$$\frac{\partial J_{\tilde{x}}(x, \bar{\sigma}, \alpha)}{\partial \tilde{x}} \Big|_{\tilde{x}=\underline{x}(\bar{\sigma}, \alpha)} = 0.$$

We have already shown that, for any x and $\sigma > 0$, $\lim_{\alpha \downarrow 0} \alpha J_{\tilde{x}}(x, \sigma, \alpha)$ is increasing in σ . It follows that, for any x and $\bar{\sigma} > 0$, there exists some $\bar{\alpha} > 0$ such that for all $\alpha \in (0, \bar{\alpha})$,

$$\frac{\partial J_{\underline{x}(\bar{\sigma}, \alpha)}(x, \sigma, \alpha)}{\partial \sigma} \Big|_{\sigma=\bar{\sigma}} > 0,$$

and therefore $\frac{\partial J_{\text{restore}}(x, \sigma, \alpha)}{\partial \sigma} \Big|_{\sigma=\bar{\sigma}} > 0$. **q.e.d.**

Appendix B: Numerical simulations in Section 3

The computations that generated the figures were conducted using a quadratic function for abatement cost and a negative natural exponential function for utility. The functional forms and parameter values used are summarized in Table B1. The flow rate μ is not given in the table; it varied as indicated in the figures and the text. The variance rate σ^2 , the ratio of the restoration cost C to the marginal cost parameter γ , and the discount rate α also vary in some figures, as indicated.

Value functions were computed via policy iteration on a “locally consistent” approximating Markov chain (see, e.g., Kushner and Dupuis, 1992). Most simulations required only 10 iterations to converge to a solution for a given set of parameter values, although more iterations were used in some cases.

variance rate		$\sigma^2 = 9.0$
discount rate		$\alpha = 0.005$
restoration cost		$C = 13000$
abatement ceiling		$\bar{a} = 20$
abatement cost	$c(a) = \phi a^2$	$\phi = 40$
utility	$u(x) = -e^{-\beta x + \kappa}$	$\beta = 0.05$
		$\kappa = -7.5$

Table B1. Parameter values and functional forms for figures.

References

- [1] Abel, A. B., and J. C. Eberly. 1994. "A Unified Model of Investment Under Uncertainty." *American Economic Review* 84: 1369-1384.
- [2] Arrow, K. J., T. Harris, and J. Marschak. 1951. "Optimal Inventory Policy." *Econometrica* 19: 250-272.
- [3] Caputo, M. R., and J. E. Wilen. 1995. "Optimal Cleanup of Hazardous Wastes." *International Economic Review* 36: 217-243.
- [4] Dasgupta, P., and G. Heal. 1974. "The Optimal Depletion of Exhaustible Resources." *The Review of Economic Studies*, 41 (Symposium): 3-28.
- [5] ———, and ———. 1979. *Economic Theory and Exhaustible Resources*. Cambridge, UK: Cambridge University Press.
- [6] ———, and Joseph Stiglitz. 1981. "Resource Depletion Under Technological Uncertainty." *Econometrica* 49: 85-104.
- [7] Dixit, A. K., and R. S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton: Princeton University Press.
- [8] Falk, I., and R. Mendelsohn. 1993. "The Economics of Controlling Stock Pollutants." *Journal of Environmental Economics and Management* 25: 76-88.
- [9] Farzin, Y. H. 1996. "Optimal Pricing of Environmental and Natural Resource Use with Stock Externalities." *Journal of Public Economics* 62: 31-57.
- [10] ———, and O. Tahvonen. 1996. "Global Carbon Cycle and the Optimal Time Path of a Carbon Tax." *Oxford Economic Papers* 48: 515-536.
- [11] Feldstein, M. S., and M. Rothschild. 1974. "Towards an Economic Theory of Replacement Investment." *Econometrica* 42(3): 393-424.
- [12] Grossman, S., and G. Laroque. 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica* 58: 25-51.
- [13] Hoel, M., and S. Kverndokk. 1996. "Depletion of Fossil Fuels and the Impacts of Global Warming." *Resource and Energy Economics* 18: 115-136.

- [14] Hotelling, H. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39: 137-175.
- [15] Kalaitzidakis, P., and S. Kalyvitis. 2004. "On the Macroeconomic Implications of Maintenance in Public Capital." *Journal of Public Economics* 88: 695-712.
- [16] Keeler, E., M. Spence, and R. Zeckhauser. 1971. "The Optimal Control of Pollution." *Journal of Economic Theory* 4: 19-34.
- [17] Keohane, N. O., B. Van Roy, and R. J. Zeckhauser. 2005. "The Optimal Management of Environmental Quality With Stock and Flow Controls." AEI-Brookings Joint Center Working Paper 05-09 (June).
- [18] Krylov, N. V. 1980. *Controlled Diffusion Processes*. New York: Springer-Verlag.
- [19] Kushner, H. J., and P. G. Dupuis. 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. New York: Springer-Verlag.
- [20] Nickell, S. 1975. "A Closer Look at Replacement Investment." *Journal of Economic Theory* 10: 54-88.
- [21] Nordhaus, W. D., and Joseph Boyer. 2000. *Warming the World: Economic Modeling of Global Warming*. Cambridge: MIT Press.
- [22] Palmieri, A., F. Shah, and A. Dinar. 2001. "Economics of Reservoir Sedimentation and Sustainable Management of Dams." *Journal of Environmental Management* 61: 149-163.
- [23] Phillips, C. V., and R. J. Zeckhauser. 1998. "Restoring Natural Resources With Destination-Driven Costs." *Journal of Environmental Economics and Management* 36: 225-42.
- [24] Plourde, C. G. 1972. "A Model of Waste Accumulation and Disposal." *Canadian Journal of Economics* 5: 119-125.
- [25] Plourde, C. G., and D. Yeung. 1989. "A Model of Industrial Pollution in a Stochastic Environment." *Journal of Environmental Economics and Management* 16: 97-105.
- [26] Rioja, F. 2003. "Filling Potholes: Macroeconomic Effects of Maintenance Versus New Investments in Public Infrastructure." *Journal of Public Economics* 87: 2281-2304.

- [27] 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.
- [28] Smith, V. L. 1972. "Dynamics of Waste Accumulation: Disposal Versus Recycling." *Quarterly Journal of Economics* 86: 600-616.
- [29] U. S. Department of Commerce (USDOC). 2004. *Statistical Abstract of the United States, 2004-2005*. Washington, D.C.: Bureau of the Census.
- [30] Ulph, A., and D. Ulph. 1994. "The Optimal Time Path of a Carbon Tax." *Oxford Economic Papers* 46: 857-868.

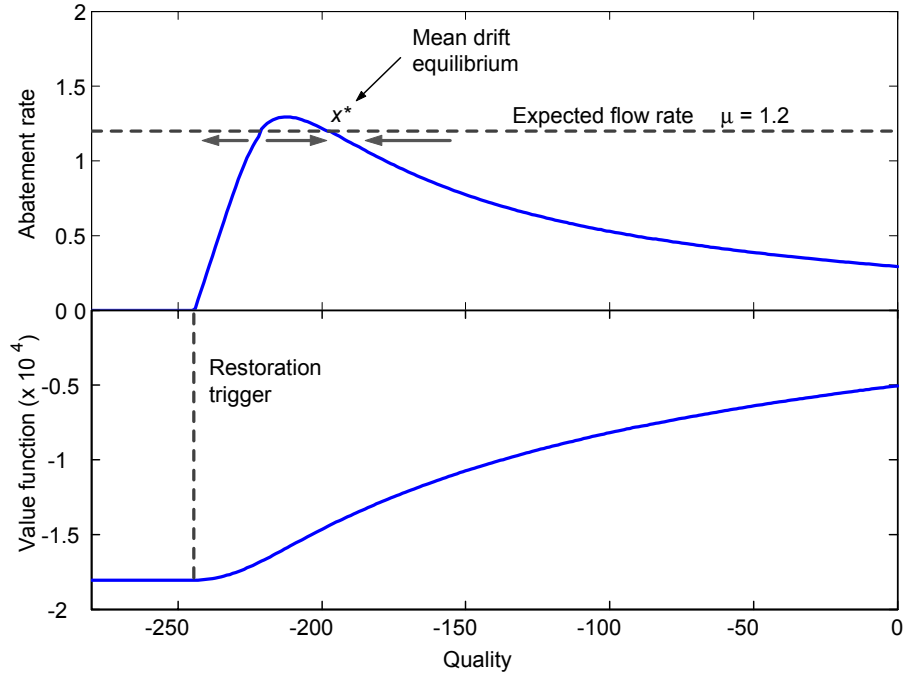


Figure 1: Optimal policies. Quality x is plotted on the horizontal axis. The optimal abatement rate, $a^*(x)$, appears above the axis, with the corresponding value function J below. Note the different units of measurement on the positive and negative segments of the vertical axis.

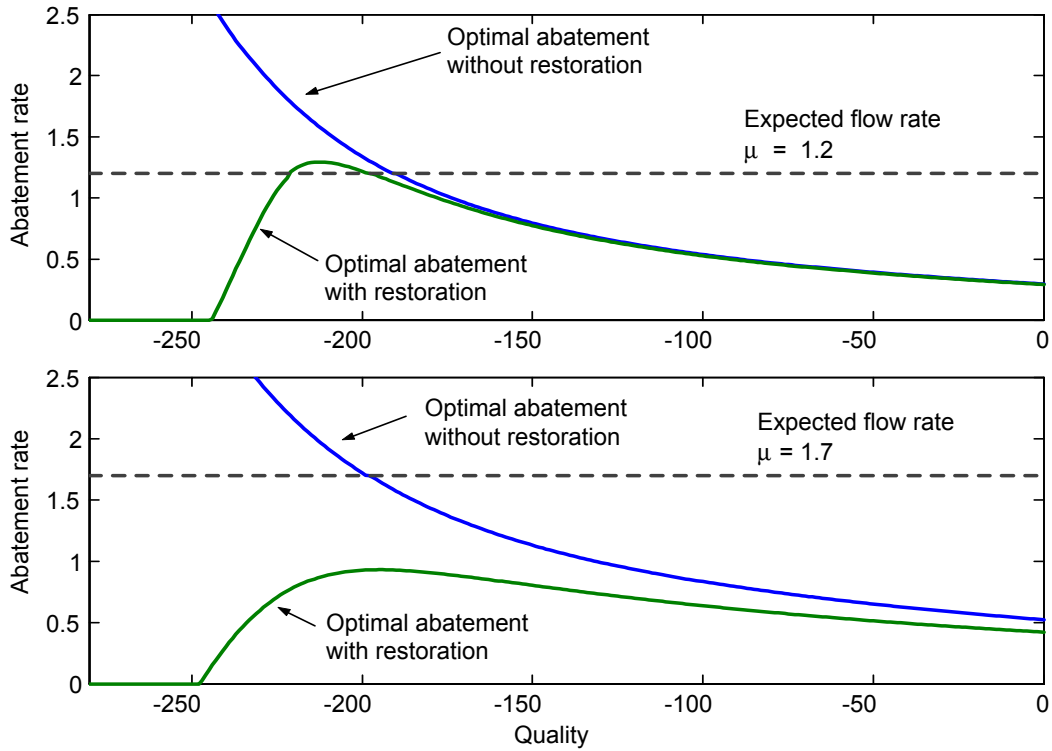


Figure 2: Effects of restoration on optimal abatement policy, for two flow rates.




Figure 3: Optimal abatement as a function of quality, for three flow rates. The left-hand figure (panel (a)) plots abatement rates; the right-hand figure (panel (b)) depicts the abatement rates as fractions of mean flow rates.

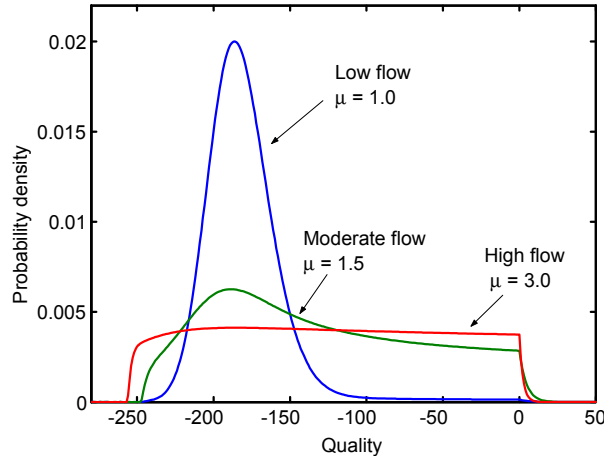


Figure 4: Frequency distributions of resource qualities (states) under optimal policies for three flow rates.

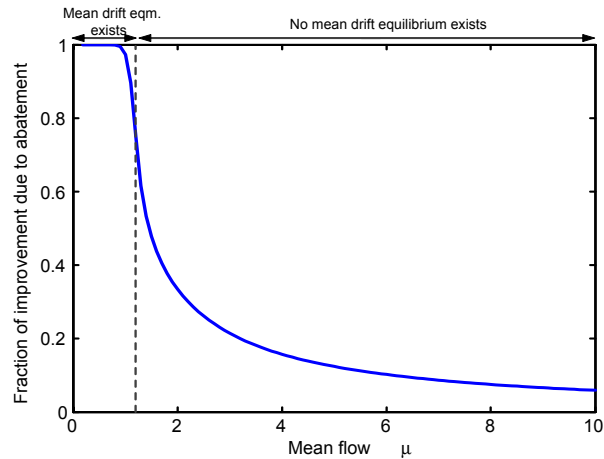


Figure 5: Fraction of total improvement due to abatement, as a function of the mean flow rate μ .

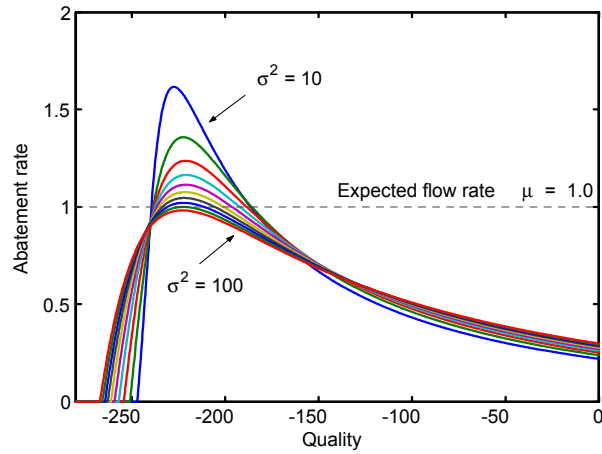


Figure 6: The effect of the variance rate σ^2 on the optimal abatement policy.




Figure 7: The fraction of quality improvement due to abatement, as a function of the relative costs of restoration and abatement. The circles mark the values of the ratio C/γ for which the fraction of total improvement achieved via abatement is 5% (hollow circles) or 95% (solid circles), for three flow rates.

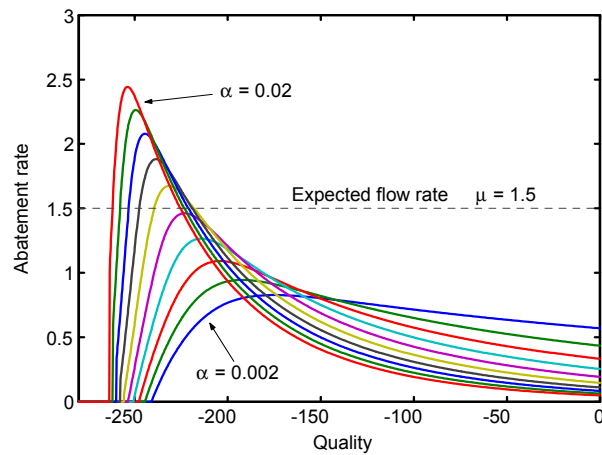


Figure 8: The effect of the discount rate α on the optimal abatement policy.