HARVARD UNIVERSITY

JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Faculty Research Working Papers Series

The Optimal Management of Environmental Quality with Stock and Flow Controls

Nathaniel Keohane, Benjamin Van Roy, and Richard Zeckhauser

June 2005

RWP05-042

The Optimal Management of Environmental Quality with Stock and Flow Controls*

Nathaniel O. Keohane
Yale University

Benjamin Van Roy
Stanford University

Richard J. Zeckhauser
Harvard University

June 2, 2005

Abstract

We characterize environmental quality as a stock, and its rate of deterioration as a flow. 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 resource stock. Abatement (curbing the flow) and restoration (restoring the stock) are interdependent tools in such problems. Under the optimal policy, periodic restoration complements positive but variable abatement that partly stems the quality decline. The preferred balance between the two strategies depends on environmental and economic factors. If flows are low enough, or if abatement is sufficiently inexpensive relative to restoration, optimal abatement may be sufficiently intense to offset the expected deterioration and produce an equilibrium in expectation. 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.

We apply the analysis to the restoration of an endangered species, and show how it could illuminate a range of other problems in the environmental arena. But the lessons are general, and we briefly discuss how they apply to the management of both physical and human capital stocks.

JEL Nos. C61, Q20, Q50.

*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, and seminar participants at Stanford University, Harvard University, and the NBER Summer Institute. Tyler Welti performed excellent research assistance.

1 Introduction

This paper considers the optimal management of a valued resource in a dynamic setting. We meld two traditional approaches to promoting quality. The first entails restoring the quality of the resource. The second involves curbing its deterioration.

Though our focus is on the control of environmental quality, a public policy problem, an example from the private sector illustrates our theory. Consider the simple problem of automobile ownership. The rate at which a car deteriorates – and hence how soon it must be replaced – can be slowed by regular maintenance. Eventually, an automobile wears out and is replaced with another: in our terms, the quality of the valued resource (transportation services) is restored. Over the life of the car, optimal maintenance varies. Early on, it increases with time: the car is still valuable but requires increasing attention. Toward the end of the car's life, however, optimal maintenance declines, since the future benefits of maintenance diminish. A driver about to junk her car is unlikely to get an oil change.

The motivating observation of this paper is that the simple logic of automobile care applies to the management of a wide range of resource stocks in the environmental arena. The accumulation and treatment of waste at landfills or generating sites is a canonical example. Optimal management both slows the generation of new wastes and periodically cleans up accumulated stocks. The management of groundwater aquifers offers another instance: an aquifer can be recharged when depleted, while conservation measures can slow the rate at which it gets drawn down. Populations of endangered species such as the California condor can be replenished through captive breeding programs; the population's rate of decline can be slowed by efforts aimed to reduce mortality risks or habitat loss. For the degraded natural habitat in the Florida Everglades, restoration would involve a large-scale project to reroute water flows and reintroduce native species, while reducing pollution would slow the rate at which degradation continues.

In each of these settings, two distinct approaches are available to manage the quality of the resource: boosting the resource stock and slowing the rate at which the stock deteriorates. Hence both *stocks* and *flows* can be controlled to promote *quality*. We refer to this class of problems as "SFQ" 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. To distinguish formally between the two strategies, 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 cleanup efforts are potentially desirable. Such scale economies are likely to obtain in many settings. For example, one way to clean up a hazardous waste site is to haul the soil away and incinerate it, in which case the costs vary little with the initial concentration of the contaminant in the soil. Similarly, there are high fixed costs involved in capping a landfill or restoring an endangered species. What is crucial to our analysis is that there are economies of scale "at the bottom" – that is, that the average costs of cleanup do not increase too much as the quality of the stock diminishes.

Under the optimal policy, the manager restores the resource whenever environmental quality falls to a sufficiently low level. At states above that point, the manager abates the new flow, at a rate that varies with the current quality of the resource. After restoration occurs, environmental degradation resumes, environmental quality starts to decline (albeit stochastically), and the cycle repeats. The optimal trade-off between abatement and restoration depends on a range of environmental and economic factors. If flows are low enough, or if abatement is sufficiently inexpensive relative to restoration, optimal abatement may be sufficiently intense to offset the expected deterioration and produce an equilibrium in expectation. Even in this case, restoration remains an option 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, when abatement is more expensive, or when restoration is less costly, the optimal policy relies more on restoration.

The model of this paper melds two instruments that have typically been considered in isolation. Conventional models of the optimal management of stock pollutants have modeled abatement alone.[1] 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] When restoration offering economies of scale is available, optimal cleanup may proceed in a jerky fashion. Rather than reaching a steady-state and remaining there, the optimal policy may entail periodic restorations punctuating long periods of abatement. Moreover,

[1]In our model, of course, the "stock" is a good (resource quality) rather than a bad (pollution); but that is only a difference in sign. The important distinction is that we consider controls on both stocks and flows.

[2]A few models of optimal cleanup of an accumulated stock of pollution 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.

the availability of restoration reduces the optimal abatement rate, with the reduction greatest as restoration becomes imminent.

In the same way, endogenous maintenance has significant consequences for optimal replacement of physical capital. In the theoretical literature on capital investment, attention has mostly centered on replacement rather than maintenance – that is, on restoration rather than abatement.[3] Optimal investment in these models typically follows an (S,s) policy (Arrow, Harris, and Marschak 1951). We show how introducing abatement changes the "trigger level" at which restoration is optimal.

The next section introduces the basic model, and formally defines our notions of abatement and restoration. Section 3 develops the theoretical results. Section 4 illustrates these results with examples of real-world SFQ problems in the environmental arena, and includes an application of the model to the recovery of the California condor. Section 5 concludes.

2 Model framework

Our model considers an environmental resource with a quality level that changes over time. In the case of accumulating waste, for example, the quality might be measured by the volume of waste: the smaller the amount, the higher the level of environmental quality.[4] In the case of an endangered species, quality would correspond to the size of the species population. We represent the quality of the environmental resource at time t by a real number x_t. Larger values of x_t represent more desirable states. We normalize the initial quality level to be equal to zero, so that $x_0 = 0$, and we shall be working mostly with negative values for x.

To keep things simple, we assume that there is a "manager" of the resource, who implements abatement and restoration policies in order to maximize the expected net present value of welfare to society as a whole. We recognize that most environmental problems do not involve a single centralized decision maker. Indeed, the manager could be the administrator of a regulatory agency that issues rules or provides rewards to influence the behavior of private-sector firms. The interactive aspects of this problem are beyond the present analysis. We focus instead on the behavior that a central planner would prescribe, recognizing full well that in many situations it would be carried out on a decentralized basis.

[3] For models of investment in physical capital, see Feldstein and Rothschild (1974) and Abel and Eberly (1994, 1996). Nickell (1975) considers maintenance as well as replacement, but he models maintenance as exogenously determined. For a model of optimal consumption of durable goods, see Grossman and Laroque (1990).

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

In addition to any efforts of the resource manager, two processes acting in opposite directions affect the level of environmental quality: ongoing damage to the resource and natural recovery processes, such as the decay of the accumulated pollutants. To capture both effects, we model the deterioration of the resource as a random variable with drift.[5] In particular, cumulative deterioration up to time t, denoted by z_t, is assumed to follow a Brownian motion with drift rate $\mu > 0$, variance rate σ^2, and $z_0 = 0$. Hence, 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, quality at time t is given by

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

Intuitively, μ can be thought of as the "average" rate of deterioration of the resource: for example, the average flow of pollution minus the natural decay of existing pollution. The random term in equation (1) captures random variations in the processes of damage and natural recovery.

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 environmental quality. Thus, at time t society derives a flow of utility $u(x_t)$ from the availability of the resource.[6] We assume that the social rate of time preference is $\alpha > 0$. We further assume that the utility function has the following properties.

Assumption 1 *The utility function u is twice continuously differentiable, with $u < 0$, $u' > 0$, $u'' < 0$, and u' unbounded above. Furthermore, $E_x \left[\int_{t=0}^{\infty} e^{-\alpha t} u(x_t) dt\right]$ 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.[7]

[5] Most of the previous literature treats pollution as deterministic. An exception is Plourde and Yeung (1989), who model pollution accumulation as stochastic and analyze efficient regulatory policies. Their analysis is akin to our model without restoration (see our Theorem 2, below).

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

[7] The assumption of negative utility is made for convenience. A reader uncomfortable with negative utility may add any constant term she wishes to make utility positive over its relevant range, without affecting the results.

We define abatement as a reduction in the rate of deterioration. Abating at rate a slows the expected deterioration rate from μ to $\mu - a$. We assume that the resource manager faces the classic trade-off between the benefits of higher environmental quality and the costs of achieving it.

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

In our SFQ framework, abatement corresponds to flow control. In the context of pollution control, for example, end-of-pipe controls on emissions and changes in the production process that reduce pollution are both forms of abatement. We assume that a finite maximum feasible rate of abatement exists, denoted $\bar{a}$.[8] This ceiling may be higher than the mean flow rate μ. Hence our model allows the resource manager not only to offset actual flows fully but to reverse the flow to some extent. In such a case, "abatement" results in a positive rate of change in the quality of the resource. Crucially, the costs of such cleanup are increasing on the margin.

Restoration corresponds to an improvement in environmental quality – affecting the stock directly, rather than by slowing deterioration. In particular, from any state x_t, we assume that the manager can restore the resource to the state $x = 0$.[9] In our analysis, we assume an extreme form of nonconvexity: a positive fixed cost of restoration to a particular level, with zero marginal cost.

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

Thus the cost of restoration is "destination-driven" in the sense of Phillips and Zeckhauser (1998): restoration cost depends on the ultimate level of environmental quality, rather than the initial level (or the amount of cleanup needed). While we assume for simplicity that costs are

[8] The assumption of a ceiling on abatement is innocuous: the ceiling can always be set high enough that the probability it binds is vanishingly small. We include it to provide 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.

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

[9] A more general framework would explicitly model the manager's choice of how much to restore – that is, how far to clean up the resource. For example, there might be more than one possible restoration method; or the costs of restoration could rise steeply as some very high level of quality was approached. But in general, such a problem will yield a single optimal destination, so little is lost by our assumption of one restoration technology. Our choice of $x = 0$ as the destination, meanwhile, is merely a normalization and does not affect the results.

destination-driven, the results of the model hold for cost functions exhibiting less extreme economies of scale, as we discuss below in section 3.4.1.

2.2 Abatement and restoration policies

We shall restrict our attention to optimal stationary policies – policies that depend solely on the state x. This assumption simplifies the analysis and has little practical impact.[10]

An *abatement policy* is a mapping $a : \Re \mapsto [0, \overline{a}]$ from the set of real numbers (the possible values of the state x) to the interval $[0, \bar{a}]$ (the feasible levels of abatement). Thus an abatement policy specifies the abatement level as a function of the state x. A *restoration policy* is characterized by a measurable closed subset R of $\Re$. Under a restoration policy R, restoration occurs whenever the state x_t occupies the set R. Given a combined abatement–restoration policy (a, R), the state of the resource evolves according to $x_t = \int_{s=0}^{t}(a(x_s) - \mu)ds + \sigma w_t - \sum_{\{i|\tau_i<t\}} x_{\tau_i}$, where τ_i is the time at which the ith restoration occurs. Hence starting from an initial state x, the infinite-horizon expected discounted utility can be written as

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

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

3 Optimal policies for SFQ problems

In this section we analyze optimal policies for general SFQ problems. We first characterize the optimal restoration and abatement policies and show how the feasibility of one clean-up method affects the optimal use of the other. We then consider how the optimal policies vary with the mean flow rate μ, the variability of the flow σ^2, the discount rate, and the costs of restoration and abatement. Finally, we briefly consider two extensions of the basic model to accommodate more general cost functions and a delay in restoration.

[10] An alternative approach would make the manager's problem one of choosing an optimal stochastic process $\{a_t\}$ measurable with respect to the filtration generated by $\{w_t\}$. However, one can show that such an optimal process can be produced by letting $a_t = a(x_t)$. Thus our assumption that the manager chooses the best among the set of stationary policies does not affect the practical implications of the analysis.

3.1 Optimal restoration and abatement

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

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

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 resource quality. It identifies two key quality levels: $\underline{x}$, the restoration trigger; and $x^\dagger$, an inflection point in the value function that coincides with maximum abatement.

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 $(\underline{x}, \infty)$. (iv) For all $x > \underline{x}$, J satisfies

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

(*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 $a^* : (\underline{x}, \infty) \longmapsto [0, \overline{a}]$ such that for every $x \in (\underline{x}, \infty)$, $a^*(x)$ uniquely attains the supremum in equation (4). (x) a^* is increasing on $\{x \in (\underline{x}, x^\dagger) | a(x) \neq \overline{a}\}$ and decreasing on $\{x \in (x^\dagger, \infty) | a(x) \neq \overline{a}\}$. (xi) Letting $R^* = (-\infty, \underline{x}]$, the pair (a^*, R^*) is an optimal policy.

This theorem (along with subsequent ones) is proven in Appendix A. Here we discuss the optimal abatement and restoration policies established by the theorem. Figure 1 illustrates optimal policies for a particular set of parameters. (The functional forms and parameter values used for all figures are provided in Appendix B.) Quality x is plotted on the horizontal axis. The optimal abatement rate, $a^*(x)$, appears above the axis, with the corresponding values of the value function J below.

Under the optimal policy, the manager restores the resource whenever environmental 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 the level of a resource's environmental quality, and a restoration corresponds to a replenishment of inventory.

As environmental quality declines from its high initial level, the optimal rate of abatement first increases, but then falls as quality approaches the "trigger level" $\underline{x}$. At a heuristic level, this policy can be understood as equating marginal benefits and marginal costs at each level of quality x.[11] From Theorem 1, the abatement rate must attain the supremum of a function $f_x(a) = aJ'(x) - c(a)$ (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.[12] 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."

Given the concave utility function, the marginal benefit from abatement at first increases as x diminishes. Since the marginal cost of abatement depends not on environmental quality, but only on the amount of the flow abated, the optimal rate of abatement rises. At a certain point, however, the optimal abatement path reverses course, with abatement *decreasing* as quality continues to worsen. Abatement reaches its peak at the inflection point $x^\dagger$ identified in Theorem 1, the point where the marginal benefit of abatement is greatest. Above $x^\dagger$, the value function is concave, so that the marginal benefit of abatement increases as the state worsens. Below $x^\dagger$, however, the value function is convex.[13] As the trigger point $\underline{x}$ nears, the marginal benefit of abatement diminishes,

[11] This and subsequent heuristic explanations aim to connect our formal results to more familiar results from economics; as a result, they are less technically rigorous than the theorems themselves.

[12] Heuristically, 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, $aJ'(x)$.

[13] 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 $x^\dagger$.

since the quality of the resource will soon be restored.[14]

Evidently, if the abatement rate rises high enough, it will equal or exceed the average flow rate μ. Suppose this condition obtains, and let x^* denote the highest quality level at which $a(x^*) = \mu$. (See Figure 1.) We call x^* an *expectation equilibrium*: since abatement at x^* equals the expected flow, the quality level will remain there in expectation.[15] Moreover, x^* is locally stable. Starting from x^*, a "high" flow of damages (at a rate greater than μ) will depress environmental quality below x^*. In response, abatement will increase, so that $a^*(x) > \mu$. Thus the quality level will return to the target x^* in expectation. A "low" flow of damages will raise quality above the target level, leading to a slackening of abatement efforts and a tendency back toward x^*.

Hence when an expectation equilibrium x^* exists, it acts as an "attractor" for the quality level, in the sense that quality will arrive at x^* and stay there in expectation. Because the optimal abatement path satisfies $c'(a(x)) = J'(x)$ at every state $x > \underline{x}$, the expectation equilibrium satisfies $J'(x^*) = c'(\mu)$: it occurs at the point where the marginal cost of fully abating expected new pollution (achieving a zero net flow in expectation) just equals the marginal net benefits from doing so. This result is analogous to the steady-state equilibrium derived in deterministic optimal control models. Nonetheless, as we show in the next section, restoration remains a remote possibility and hence affects the optimal abatement path.

Because a declines to zero as the trigger level approaches, there must be a second state $x^{**} < x^\dagger < x^*$ at which optimal abatement again equals the average flow rate. This too is an equilibrium, but it is unstable: any deviation in flow from μ will tend to push quality either downward toward $\underline{x}$ or upward toward x^*. If flows are large enough to push quality below x^{**}, the state will (again in expectation) decline to the restoration trigger $\underline{x}$.

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 availability of abatement changes

[14] Note that the abatement rate falls to zero at the restoration trigger point $\underline{x}$. The marginal benefits of further abatement at that point are zero, because the state will be restored to $x = 0$ immediately. The smooth-pasting condition (Krylov, 1980) requires that as the state approaches the trigger, the marginal benefits from abatement decline smoothly to zero. Marginal cost must follow suit, implying that abatement must go to zero as well.

This result would change slightly under the more general cost functions considered in section 3.4.1. Given a variable component of restoration cost $\gamma(x)$,with $\gamma'(x) < 0$, the smooth-pasting result would still hold; but abatement would decline smoothly to $\underline{a} > 0$ satisfying $c'(\underline{a}) = -\gamma'(\underline{x})$.

[15] We use the term "equilibrium" in the sense of "system stability," rather than in a strategic or game-theoretic sense. The "expectation" refers to the frequency distribution of states, rather than the beliefs of economic agents.

the quality level that triggers restoration of the resource, while the feasibility of restoration reduces the optimal rate of abatement.

To see how the possibility of abatement affects optimal restoration, consider the case in which only restoration is feasible. Because the absence of abatement as an option constrains the manager's actions, the value function in the restore-only case must be everywhere less than the value function in the combined abate-and-restore case. In other words, the availability of abatement makes the resource more valuable, because its deterioration can be slowed. Figure 2 illustrates this result.

While introducing abatement raises the value function at every state, the magnitude of the shift is not the same everywhere. Thus the feasibility of abatement changes the slope of the value function. An immediate implication is that introducing abatement may change the optimal restoration trigger. To see why, recall that restoration occurs at the state where the value function equals $J(0) - C$. Consider what happens as quality worsens from $x = 0$. The steeper is the value function below $x = 0$, the sooner it will have declined by enough to make restoration worthwhile.

Figure 2 presents a case in which introducing abatement raises the optimal restoration trigger. At first glance this situation may seem counterintuitive: why are not restoration and abatement always substitutes for one another? The reason is that introducing abatement may raise the value function more at high levels of quality than at low levels. When quality is high, restoration is distant, and abatement allows the manager to maintain quality at a higher level than would be possible in its absence. Hence introducing abatement may raise the value function more at high levels of quality than at low levels. If so, the restoration trigger must rise. Different circumstances, however, can produce the opposite result. Simulations demonstrate that the availability of abatement lowers the restoration trigger when mean flow is sufficiently low. The precise conditions under which the trigger falls or rises provide an avenue for future research.

In contrast, the potential for restoration always lowers abatement. To see this, consider the other polar case, in which only abatement is possible. Let J_{abate} denote the optimal value function in this case:

$$J_{\text{abate}}(x) = \sup_a E_x^a \left[\int_{t=0}^{\infty} e^{-\alpha t} (u(x_t) - c(a(x_t))) dt \right], \tag{5}$$

where the supremum is taken over abatement policies. Let a_{abate} be the corresponding optimal abatement policy. Using similar methods as in the proof of Theorem 1, one can show that $J_{\text{abate}} < 0$, that $J_{\text{abate}}(x)$ is finite for every x, and that J_{abate} is twice continuously differentiable and concave with J'_{abate} unbounded above; moreover, an optimal policy a_{abate} exists and uniquely satisfies the corresponding Bellman equation. Theorem 2 describes the optimal abatement policy in this case. Abatement rises monotonically as quality diminishes, up to the abatement ceiling $\bar{a}$. One consequence is that an expectation equilibrium *always* exists in the abatement-only case, as long as the abatement ceiling is higher than the mean flow rate.

Theorem 2 *Let Assumptions 1 and 2 hold, and assume that restoration is not feasible. Then (i) there exists a state* $\hat{x}$ *such that* a_{abate} *is decreasing on* $(\hat{x}, \infty)$ *and* $a_{abate}(x) = \overline{a}$ *for all* $x \leq \hat{x}$*; (ii)* $\lim_{x\to\infty} a_{abate}(x) = 0$*; and (iii) if* $\bar{a} > \mu$*, then there exists a state* x^*_{abate} *such that* $\mu < a_{abate}(x)$ *for* $x < x^*_{abate}$ *and* $\mu > a_{abate}(x)$ *for* $x > x^*_{abate}$*.*

Next, Theorem 3 compares the optimal value functions and abatement policies in the abatement-only case with those in the full SFQ case.

Theorem 3 *Let Assumptions 1, 2, and 3 hold. Then (i)* $J' < J'_{abate}$*; and (ii) for each state* $x \in (\underline{x}, \infty)$*, where* $\underline{x}$ *is the trigger point under restoration, either* $a(x) < a_{abate}(x)$ *or* $a(x) = a_{abate}(x) = \overline{a}$*.*

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

Assertion (ii) establishes that the optimal abatement rate when restoration is available is less than in the abate-only case (unless, of course, the abatement ceiling binds). When restoration is possible, the present value of net benefits (i.e., $J(x)$) increases more slowly as the state improves. Loosely speaking, the marginal gains from abatement are lower. Hence less abatement is optimal. A corollary is that even if an expectation equilibrium exists when restoration is feasible, it must occur at a lower level of quality. That is, $x^* < x^*_{\text{abate}}$. Thus the availability of restoration alters the optimal abatement policy, even when the probability of its occurrence is very low.

Figure 3 summarizes this discussion, portraying optimal abatement policies with and without the possibility of restoration. In the top panel, an expectation equilibrium exists when restoration is available, but it occurs at a lower level quality than if restoration were not an option. The bottom panel of Figure 3 illustrates a case in which restoration has a more drastic consequence for the abatement path. No expectation 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 until the trigger level is reached, and then restores the resource. Which of the two cases portrayed in Figure 3 prevails depends on how fast the resource deteriorates, as we discuss in the next section.[16]

3.3 The role of economic and environmental factors

In this section, we consider how the optimal policies just described vary with parameters of the physical environment (in particular, the mean and variability of the rate of deterioration), the nature of preferences (including the discount rate), and the costs of cleanup. We rely for the most part on simulations for our results, since important relationships in the model are often complex, hence resistant to straight analytic demonstrations. Where analytic results can be established, we state them. We discuss a range of salient features of the model in turn: the value of the resource; the abatement path, including the existence of an expectation equilibrium; the "trigger" state at which restoration occurs; and the balance between restoration and abatement.

3.3.1 The value of the resource

It is straightforward to show that the higher are flows, the lower is the value of the resource at every quality level. Higher flows mean greater deterioration in the resource, and hence bring both lower utility and higher cleanup costs. The following theorem summarizes this point. The proof is straightforward and hence omitted.

Theorem 4 *Let Assumptions 1, 2, and 3 hold. Let $J(\cdot, \mu)$ be the optimal value function given a drift rate $\mu > 0$. Then, for any x, $J(x, \mu)$ is decreasing in μ.*

The effect of the variance rate σ^2 on the value function turns out to be ambiguous in the full SFQ case. Thus it is useful to consider its role when only one of the cleanup options is feasible. When only abatement is feasible, the value function J_{abate} is decreasing in σ^2. In that case, the value function is concave everywhere: hence by the usual Jensen's Inequality arguments an increase

[16] A third possibility exists: the optimal abatement rate may equal average flow at exactly one point, so that the abatement function is tangent to the horizontal line at μ. The resulting "expectation equilibrium" x^* is then stable from the right but not the left. For $x > x^*$, abatement will be lower than μ and the state will tend to return to x^*. For $x < x^*$ abatement will also be less than μ, implying that in expectation the quality level will decline to the trigger level for restoration.

in variance lowers the value of the resource in expectation. This result is stated in the following theorem, the proof of which is omitted.

Theorem 5 *Let Assumptions 1 and 2 hold. Let $J_{abate}(\cdot, \mu, \sigma^2)$ be the optimal value function given the variance rate σ^2, when only abatement is feasible. Then, for any x, $J_{abate}(x, \mu, \sigma^2)$ is decreasing in σ^2.*

In the restoration-only case, the opposite result holds, as long as the resource manager is sufficiently patient. That is, J_{restore} is increasing in σ^2 for small enough values of the discount rate α. In this case, the value function is convex near the restoration trigger, so the Jensen's Inequality argument no longer holds. Rather, the intuition behind the restoration-only result is as follows: When flows are more variable, we are more likely to have recently restored the resource at any point in time, since the restoration threshold is more likely to have been crossed. Hence the greater the variance, the more likely is the resource to be at a high level of quality. If the discount rate is sufficiently low, this exchange of more numerous cleanups for higher quality makes variability a welcome companion: it raises the present value of expected utility.

Theorem 6 *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 $\overline{\alpha} > 0$ such that for any $\alpha \in (0, \overline{\alpha})$, $J_{restore}(x, \sigma^2, \alpha)$ is increasing in σ^2.*

Finally, the value of the resource declines with the costs of abatement and restoration, as must be the case.

3.3.2 The abatement path

Figure 4 illustrates optimal abatement policies for three values of the mean flow rate μ. Note that the vertical axis measures the fraction of mean flow abated, so that the abatement rates are normalized by the flows. When the mean flow rate 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 in the environment for a shorter period of time before being cleaned up. Hence at higher flow rates, restoration becomes more attractive relative to abatement, and less abatement is done.

Figure 5 depicts the effects of the variance rate σ^2 on the optimal abatement path. Note that abatement reaches a higher maximum when the variability is lower. This is a consequence of

Jensen's Inequality. Because abatement costs are convex, abatement becomes more expensive in expectation as the variability rises.

The effects of economic variables accord with intuition. When marginal abatement cost is higher, the optimal abatement rate is lower at every state. Figure 6 illustrates this point, plotting optimal abatement for a range of marginal cost functions (the parameter γ corresponds to the slope of the linear marginal abatement cost function used in the simulations). The effect of higher discount rates, depicted in Figure 7, is perhaps more interesting. When the quality of the state is high, a higher discount rate leads to less abatement – as might be expected. As the state deteriorates, however, this relationship inverts: abatement becomes greater at higher discount rates, and thus reaches a higher peak. Putting more weight on current relative to future utility results in more abatement, because the benefits from eventually restoring the resource are discounted more heavily. This result is intuitive, and yet leads to a surprising conclusion. In conventional models of the management of resource stocks, greater patience leads to *more* abatement in the near term, not less. In the SFQ model, over a range of states a more patient manager will be more tolerant of environmental degradation in the short run – because such tolerance will hasten restoration, and raise environmental quality in the long run.

An important implication of these results is that the existence of an expectation equilibrium, and thus the relevance of restoration, depends on the characteristics of the environment and of the economic system. When flows are high, or highly variable, an expectation equilibrium may not exist in the full SFQ case. On the other hand, an expectation equilibrium always exists in the abatement-only case, as we showed above. Similarly, when abatement is expensive or the discount rate is low, an expectation equilibrium is less likely to exist when restoration is feasible. The availability of restoration matters most for management in such cases.

3.3.3 The restoration trigger

The restoration trigger for each abatement path is the point at which abatement falls to zero. In Figure 5, for example, the restoration trigger is lower when flows are more variable. In Figure 4, on the other hand, the effect of the flow rate on the restoration trigger is ambiguous.

Figure 8 plots the restoration trigger against the flow rate μ and the variance rate σ^2, for a particular range of parameter values. The top two panels represent two-dimensional projections, showing how the trigger varies with one parameter or the other. (Note that μ *decreases* from left to right in the top left panel, for consistency with the view in the bottom graph.) The lower panel shows the three-dimensional surface over the same range. The restoration trigger is measured on the vertical axis; notice that μ increases "into" the page, while σ^2 increases from left to right.

When flows are sufficiently high, the restoration threshold is inversely related to both the flow rate and its variability. That is, higher flows and greater variability both drive the restoration trigger down. For example, when $\mu = 5$ the trigger point $\underline{x}$ declines smoothly from -263 to just under -270 as σ^2 increases from 1 to 100. We offer a heuristic explanation for this result, which borrows from the theory of real options.[17] 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 the cost of restoration is unaffected by quality. The resulting option value represents a reward to waiting, and so leads to a lower restoration trigger. Moreover, the reward is greater when flows are more variable.

This result does not hold at all flow levels, however. In particular, for low flow rates (not shown in Figure 8) an increase in variability raises the restoration trigger. We conjecture the following: if an expectation equilibrium does not exist, the restoration trigger decreases monotonically with the variability of flows.[18] When an expectation equilibrium exists (i.e., if flows are low enough), it interrupts the decline of quality towards the restoration trigger, and as a result may alter the relationship between flows and restoration.

The effect of the flow rate μ on the restoration trigger appears to follow a similar pattern. In Figure 8, the trigger rises as mean flow falls, down to $\mu = 1.5$. Below that point, the trigger falls with further decline in the flow rate; indeed, the drop is precipitous from $\mu = 1.0$ to $\mu = 0.5$ (not shown on the figure). We conjecture that whether the restoration trigger increases or decreases in the flow rate depends on whether an expectation equilibrium exists. The relationship between abatement cost and the restoration trigger is similarly ambiguous. However, the restoration trigger does appear to decline monotonically with the discount rate. That is, the more weight is put on current rather than future welfare, the further the resource is allowed to deteriorate before restoration occurs.

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

[18] Our conjecture is a statement of sufficiency, not necessity. Simulations show that the trigger falls with variability even in some cases with expectation equilibria. Indeed, for the lowest flow depicted in Figure 8 ($\mu = 1.0$), an expectation equilibrium exists for almost all the entire range of σ^2.

The effect of restoration cost is unambiguous, and is stated in Theorem 7. As the cost of restoration increases, the trigger level $\underline{x}$ decreases. As intuition would suggest, a higher cost increases the incentive to delay restoration.

Theorem 7 *Let $\underline{x}(C)$ denote the trigger level for a given restoration cost $C > 0$. Then $\underline{x}(C)$ is decreasing in C.*

3.3.4 The optimal balance between restoration and abatement

So far we have focused on the particulars of the optimal policy, e.g., the peak of the abatement path or the location of the restoration trigger. Here, we take a broader view, summarizing the effects of economic and environmental variables on the balance between abatement and restoration.

Figure 9 shows how the importance of abatement relative to restoration varies with the flow rate. The horizontal axis measures the flow rate. The vertical axis measures the time-averaged rate of abatement as a fraction of the flow of damages, or (equivalently) the fraction of total damages that is cleaned up optimally by abatement rather than restoration. Thus for a given flow rate, the height of the curve represents the fraction of cleanup due to abatement. The remainder of the cleanup, from the curve to the top of the graph, is due to restoration. For example, at a flow rate of 3, approximately 20% of the total flow of damages is cleaned up by abatement, with the remaining 80% cleaned up through periodic restorations.

Figure 9 also illustrates how the existence of an expectation equilibrium depends on the flow rate. The dashed line on the figure marks the critical value of the flow rate that determines whether or not an expectation equilibrium exists. (In Figure 9, this critical value is just above 1.2.) For flow rates below this critical value, abatement offsets expected flows completely over some range, so that an expectation equilibrium is reached. In this case restoration occurs with very small probability, and virtually all of the damages are cleaned up through abatement. For low flow rates, then, we may say that abatement is the "principal strategy." When flows increase beyond this cutoff, the expectation equilibrium vanishes, and restoration becomes the principal strategy.

Similarly, restoration predominates as flows become more variable, or as abatement becomes more costly. Plotting the fraction of cleanup by abatement against the variance rate σ^2 or the marginal cost parameter γ would reveal a similar pattern as in Figure 9. Since greater variability of flows depresses the optimal abatement rate, the share of cleanup achieved by abatement falls as the variability rises. Similarly, as γ increases, so that abatement cost rises more rapidly, the fraction of cleanup achieved through abatement falls. On the other hand, an analogous plot of cleanup against the discount rate would yield the reverse relationship, with abatement accounting for more cleanup as the discount rate increases.

A particularly interesting feature of the SFQ model is that the optimal balance between restoration and abatement is highly sensitive to small changes in key parameters. Consider the effects of changes in the flow rate around the critical value of 1.2 in Figure 9. As the mean flow rate increases from 1.1 to 1.3, the fraction of cleanup produced by abatement drops dramatically, from 0.9 to 0.6. Between $\mu = 1$ and $\mu = 2$, the abatement fraction falls from 0.97 to 0.33. The optimal mix of cleanup methods is similarly sensitive to changes in variability, marginal abatement cost, and the discount rate around the critical values that determine whether an expectation equilibrium exists.

Whether cleanup relies more on restoration or on abatement determines how the quality of the resource varies over time. Figure 10 plots the frequency distributions of states for three flow rates. When flows are low, an expectation equilibrium is achieved. States close to this equilibrium level are much more common than other states. Indeed, the peak of the frequency distribution for $\mu = 1$ occurs precisely at the expectation equilibrium ($x = -185$) depicted in Figure 4 for the same flow rate. At somewhat higher flow rates, no expectation equilibrium exists, and restorations occur more frequently. As a result, high-quality states become relatively more common, flattening the frequency distribution. For the moderate flow rate depicted in the figure ($\mu = 1.5$), the distribution retains a peak, occurring just above the point at which abatement reaches its maximum. At the high flow rate, restoration becomes even more important relative to abatement, and all states between the initial quality level and the restoration point occur with roughly equal frequency.[19]

3.4 Extensions

3.4.1 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 exhibiting less extreme economies of scale. For example, suppose that the cost of restoring the resource to state 0 starting from quality level x has a fixed component F, as before, but also has a variable

[19]Figure 10 also demonstrates a seemingly paradoxical result: in the full SFQ case, when both restoration and abatement are possible, the average quality of the resource may sometimes be higher when flows are greater. When the flow of damages is high, restorations may become sufficiently more frequent that high-quality states are more common.

component $\gamma(x)$. We would expect $\gamma(x)$ to be a decreasing function of x; i.e., the restoration cost increases with the amount of restoration done. 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. If we now 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.

If $\gamma(x)$ 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.

3.4.2 Delayed restoration

In many real-world applications of the SFQ model, restoration is unlikely to be instantaneous. For example, restoring a degraded habitat, such as the Everglades, may take several years. Similarly, proposed methods to remove greenhouse gases from the atmosphere (e.g., seeding oceans with iron filings to promote the growth of carbon-dioxide-absorbing plankton) would require long lead times.

Consider a generalization of the model formulated in Section 2, where a delay of length D is incurred in restoration. During the interval $[\tau_i + D, \tau_{i+1})$ between the completion of the ith restoration project and the commencement of the next, the state evolves according to $x_t = \int_{s=\tau_i+D}^{t} \left((a(x_s) - \mu)ds + \sigma dw_s\right)$.The optimal value function is now defined by

$$\begin{aligned} J(x) &= \sup_{a,R} E_x^{a,R}\left[\int_{t=0}^{\tau_1} e^{-\alpha t}(u(x_t) - c(a(x_t)))dt \right. \\ &\left. + \sum_{i=1}^{\infty} \int_{t=\tau_i+D}^{\tau_{i+1}} e^{-\alpha t}(u(x_t) - c(a(x_t)))dt - \sum_{i=1}^{\infty} e^{-\alpha \tau_i} C\right] \end{aligned} \tag{6}$$

(compare to equation (3)). Note that states observed during a restoration project do not enter into equation (6). Instead, the restoration cost C now 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).

Our analysis can be extended to establish a generalization of Theorem 1. The only difference is that the boundary condition stated in the first part becomes $J(x) = e^{-\alpha D} J(0) - C$ for all $x \leq \underline{x}$. The other results in this section continue to hold as well (in particular, Theorems 3 and 4.)[20]

[20] A slight modification is needed to Theorem 7: the trigger level $\underline{x}$ is nonincreasing (rather than decreasing) in C.

4 SFQ problems in the environmental arena

In this section, we explore the implications of our theoretical model for the management of resource stocks in the real world. First, we briefly sketch the model's application to a number of environmental problems, ranging from accumulating waste to zebra mussels. We then consider the case of the California condor in detail.

4.1 Environmental applications

The accumulation of wastes at disposal sites or generating facilities is a canonical SFQ problem. Consider the optimal management of municipal solid waste, for example. The environmental quality of a landfill site and the surrounding area diminishes as solid waste 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.[21] In a typical scenario, waste diversion remains roughly constant over time, or changes only with changing preferences (i.e., a desire to increase levels of recycling) or prices (e.g., land becomes more expensive, or recycled materials become more valuable). Optimal waste management, on the other hand, would vary the rate of abatement over time. When a landfill is first opened, diversion should be relatively high. That is because waste dumped early will be around for nearly the landfill's entire life. Thus, the discounted expected damages it imposes will be high relative to the damages from garbage arriving later. As the landfill nears capacity, waste diversion should slow, since the waste will impose damages only until the time of restoration.

Similar issues, on a different time scale, are involved in the management of hazardous wastes. Consider the chemistry department at Harvard University.[22] The department's laboratories accumulate a variety of toxic and reactive substances. Storing such substances on campus heightens health and fire hazards.[23] Removing the wastes for permanent disposal – "restoration," in this

In this model of delay, if the restoration cost C is sufficiently low, the optimal policy will be to continually restore the resource – effectively avoiding the negative utility that comes from letting the resource deteriorate.

[21]With 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 the exposure costs and restoration costs rise as we move to successively more expensive landfills. Essentially, this would append results from the theory of nonrenewable resources to our models (Dasgupta and Heal, 1979; Hotelling, 1931). Abatement today would be influenced by the shadow price of future restorations.

[22]We thank Henry Littleboy, Health and Safety Officer (for Harvard's Faculty of Arts and Sciences Office of Enviornmental 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.

[23]Of course, chemical waste storage and disposal are heavily regulated by the Environmental Protection Agency.

context – involves economies of scale, reflecting the fixed costs of labor and transportation. Chemical wastes are hauled away in "lab packs": containers are collected from labs and packed in larger drums with wastes of similar types. A 55-gallon drum of corrosive flammable liquids costs $320 to ship; a 30-gallon drum costs $215, and a single 5-gallon container $95.

At least in principle, several methods exist to control the flow of lab waste generated: experiments could be curtailed or altered 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, albeit at significant cost. Although some abatement would likely be optimal, little concerted effort is actually made to curb flows. A partial explanation is that until recently, individual laboratories were not charged for disposal, and thus had little incentive to reduce their chemical use. Individual labs recently began to pay a volume-based charge for both solvents and lab packs. Limited experience indicates that the use of chemical wastes is fairly inelastic, suggesting high costs of substantial abatement.

A very different application is the sedimentation of reservoirs.[24] The "stock" in this context is the capacity of the reservoir, which is diminished as sediment flows into the reservoir and accumulates. Dams are commonly designed to have finite lives: the reservoir behind a dam fills up with sediment, until the dam is retired and a replacement dam is constructed. Retirement and replacement constitute an extreme form of restoration, and one whose costs are essentially destination-driven. Alternatively, the stock of sediment can be removed directly by siphoning or dredging – activities likely to exhibit economies of scale. Meanwhile, a range of strategies exists to abate the sediment flow. The flow into the reservoir can be reduced by soil conservation, reforestation, and other measures in the catchment area; or sediment can be routed away from the reservoir. The precise nature of the optimal policy will depend on site-specific costs of stock and flow controls. Nonetheless, the common practice of letting sediment accumulate unchecked before retiring the dam – equivalent to a restoration-only policy – is almost surely suboptimal. At the same time, "sustainable management" that seeks to maintain an equilibrium by relying exclusively on sediment flow control, without periodic restoration (Palmieri, Shah, and Dinar 2001), is equally unlikely to be optimal.

The SFQ model also applies to the control of animal pests, such as zebra mussels (*Dreissena*

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.

[24] The information about dam sedimentation is taken from Palmieri, Shah, and Dinar (2001).

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 area of the United States has been estimated to have cost as much as $1 billion in the 1990s alone.[25]

The feasibility of preventing mussel settlement varies by location, and control strategies vary accordingly. 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 (thus shutting down the plant). An (S,s) policy is followed. Mussels are allowed to settle and grow, and periodically are removed. Removal is done every year or two in western Lake Erie, their densest habitat. Inside the plumbing systems of power plants and water works, mussels are inaccessible to mechanical removal, but chemical removal is feasible. In such locations, both flow and stock controls are employed. Continuous low-level chlorination of circulating water is an abatement policy. It kills juvenile mussels, and inhibits their settlement. Periodic (annual or semi-annual) injections of high concentrations of chlorine represent a restoration strategy used to kill off adult mussels that have settled.[26]

Our model also applies well to the degradation of natural habitats such as the Florida Everglades, where natural features and both healthy and endangered species comprise the stock of environmental quality. The area's deterioration can be slowed by pollution prevention efforts and by reductions in agricultural and urban water use. Environmental quality can be restored by recreating historical water flows and reintroducing native species.[27] Of course, rerouting freshwater flows

[25]Personal communication, Charles O'Neill, Project Director, National Zebra Mussel Information Clearinghouse, New York Sea Grant.

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

[27]While we have modeled a cycle of cleanup and deterioration, in some real world cases (such as the Everglades) restoration is likely to be done only once, after which the process of environmental degradation stops. Nonetheless, the implications of the model are reasonably clear for such cases: starting from an already degraded state, we abate at the margin (if possible) and then eventually restore all at once. In order for one-time restoration to be optimal, however, there would presumably need to be a nonconvexity in benefits, or at least very high marginal damages from the first few units of degradation.

Note also that while many real-world restorations are treated as "one-time-only" events, they are better modeled as combined abatement-restoration policies in which an expectation equilibrium exists at a fairly high level of environmental quality – implying that the probability that the resource quality declines far enough to trigger restoration is very small, but not zero. For example, a manager who cleans up a river may not expect to do so again. The to the Everglades will require large construction projects, and the effects will not be immediate. Thus the appropriate SFQ model incorporates delay. The same point holds more generally for the restoration of natural habitats or endangered species populations, such as the condor example discussed in the next section. Nonetheless, the core methodological lessons of the model apply, as we showed in Section 3.4.2 above.

4.2 California condor restoration

One of the most prominent efforts at endangered species restoration has been the successful recovery of the California condor (*Gymnogyps californianus*). With a wingspan of nine and a half feet, the condor is the largest bird in North America. Until the mid-nineteenth century the condor's range extended as far north as the Columbia River Gorge and south into Baja California.[28] Throughout the twentieth century the population declined precipitously; in crisis years, it fell from approximately 100 birds in the 1940s to 22 in 1982. The decline appears to have been due to reduced reproduction (perhaps a result of DDT) and to human-created mortality, including lead poisoning from bullets in game carcasses, hazards from man-made structures such as powerlines, and the shooting of the condors themselves.

The condor was recognized as "endangered" by the Federal government in 1967. The ensuing recovery plan, approved in 1979, recommended captive breeding of condors to aid recovery. After a disastrous winter of 1984-85, when six of the remaining fifteen wild condors died, the U. S. Fish and Wildlife Service decided to capture all of the remaining wild birds for the captive breeding program. The reintroduction program was launched in 1992 with the release of two captivity-born juveniles into the Sespe Condor Sanctuary in Los Padres National Forest. By October of 2003, the wild population had climbed to 83 birds, including one chick hatched in the wild.

The captive breeding program exemplifies restoration with economies of scale. The costs of captive breeding are largely fixed: the costs of developing a program to capture the birds, maintaining the condor populations in captivity, and releasing condors back into the wild do not change greatly on the margin with the number of birds released. Abatement measures include the provision of food carcasses such as stillborn calves (to prevent lead exposure); promotion of alternatives to lead in ammunition; prohibitions on shooting the condors; protection of suitable habitats; and attempts

possibility remains – however slight – that unforeseen changes in water currents or waste generation will create the need for future restorations.

[28]The diaries of Meriwether Lewis and William Clark report several sightings of the "Buzzard of the Columbia" in 1805 and 1806.

to limit injuries from powerlines and other human structures. Each of these measures potentially helps slow or offset some portion of the decline in the condor population. Arranging them in order of increasing cost per condor saved, as efficient policy requires, produces an increasing marginal abatement cost curve.

Simulating the optimal management policy for the condor requires estimating the restoration cost, abatement cost, and utility functions. We estimated the costs of restoration and abatement using detailed program cost estimates in the Recovery Plan (U.S. Fish and Wildlife Service 1996). Details of how this was done can be found in Appendix C. The expected flow rate μ was set to 2.9 condors per year, equal to the average annual rate of decline of the condor population over the decades prior to restoration, based on wild condor censuses.

Because no studies exist of willingness to pay for condor populations, we took point estimates from two studies of the value of bald eagle populations (Boyle and Bishop 1987; Stevens et al., 1991). We divided each estimate by 10 to yield two more estimates, and then multiplied by the current U. S. Sierra Club membership of 700,000. We thus arrived at four point estimates of total annual willingness to pay ranging from \$1.1 million to \$23 million. A concave utility function was fitted to each point estimate, assuming that the willingness to pay for a wild population of zero birds is zero.[29] These estimates are deliberately conservative. The low estimate assumes that Sierra Club members – and only Sierra Club members – would each be willing to pay only \$1.53 annually to preserve a population of three hundred birds, and would pay a *maximum* of \$1.75 for an arbitrarily large population. In the highest-valuation case, annual per-capita willingness to pay would never exceed \$37, again for an arbitrarily large number of condors.[30]

Figure 11 illustrates simulated optimal abatement polices for these four estimates of willingness to pay. At the lowest valuation, the simulated optimal abatement path is essentially tangent to the horizontal line at the mean flow rate of 2.9, with the peak occurring at a population of 138 condors. (Abatement rates, like flows, are measured in condors per year.) At higher valuations, optimal abatement exceeds mean flow over a significant range. The resulting "expectation equilibria" occur at populations of between two and three hundred condors.

Of greatest interest are the optimal trigger levels for restoration suggested by the model. Recall

[29]We used a utility function that exhibits much more rapidly diminishing marginal utility for condor populations than some results in the literature on endangered species valuations have suggested. For example, a meta-analysis by Loomis and White (1996) suggests a utility function of the form ax^b with $b = 0.8$ – a function that results in WTP estimates for large populations that are orders of magnitude greater than ours.

[30]The cost of a regular Sierra Club membership in 2003 was \$39.

that the trigger levels represent the quality level of the resource at which the gain in the value function just exactly equals the cost of restoration. In this context, therefore, restoration is optimal when the present value (discounted expected utility) of the wild condor population falls $15 million below the value associated with the population once restored. The trigger levels in our simulations correspond to populations of 85, 149, 189, and 205 condors. For comparison, the actual condor population in the wild fell to 15 birds before restoration was undertaken.

To explore the sensitivity of our conclusions to the parameter assumptions, we ran Monte Carlo simulations varying three parameters: per-capita willingness to pay, an abatement cost parameter, and the variability of the flow rate. (See Appendix C for details.) The central policy in the simulations reaches a maximum abatement rate of 3.4 condors per year, with an average of 1.1. That policy exhibits an expectation equilibrium at a population of 246 condors. The restoration triggers in the 95% confidence region vary from 92 to 193 birds, with a mean of 164.

Hence over a wide range of parameters, our model suggests that the condor population was allowed to fall much too far before restoration was undertaken. Indeed, a farsighted policy might have resulted in a expectation equilibrium population of a few hundred birds, resorting to restoration only if several bad shocks reduced the population far enough.

5 Conclusion

In a wide range of settings, both stocks and flows can be controlled to improve the quality of a 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 characteristics of the physical environment, the nature of preferences (in particular the discount rate), and the costs of cleanup. If flows are low enough or abatement is sufficiently inexpensive, an "expectation equilibrium" may be reached where abatement efforts just offset the expected deterioration of the resource. In that case, abatement is the principal cleanup strategy, although the optimal abatement rate is still lowered by the potential for restoration. 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 in the environmental arena: the disposal of municipal solid waste and hazardous laboratory waste; the prevention of siltation in reservoirs; the control of pests such as zebra mussels; the regeneration of natural habitats such as the Florida Everglades; and the recovery of the endangered California condor. Note that these problems involve private decisions (lab wastes), public decisions (municipal solid waste), and situations where the public and private sectors interact, often in a regulatory context.

The analysis generalizes readily to the management of physical and human capital. Optimal policies for replacing a machine, investing in capital equipment, or purchasing consumer durables cannot be derived independently of the optimal maintenance paths. Indeed, how far the productivity of capital should be allowed to fall before replacement depends not only on how costly is replacement but also on how rapid and variable is deterioration. 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 lay off its older workers, or reassign them to tasks where the latest technical skills are less essential and replace them with recently trained workers. In this context training is "abatement" and replacement – which incurs costs such as severance payments or raised experience-rated unemployment insurance – amounts to restoration.

Government will likely play a role in many SFQ problems, e.g., in controlling environmental quality, and must understand the central lesson: 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) will be more expensive than achieving the same present value of expected utility from quality, but allowing quality to vary over time. A policy relying solely on restoration will not only restore too frequently (since deterioration is unchecked), but will also allow quality to fall too far before each restoration (since the optimal trigger rises when abatement is available).

These errors are likely to be relevant to real-world environmental policies that utilize only one cleanup strategy or the other. For example, endangered species laws in the United States allow animal populations to fall to critical levels before government intervenes. Costly and risky restorations, such as the condor restoration currently underway, are one consequence. A better policy would seek to stem population decline much earlier, for example by giving landowners incentives to maintain appropriate habitat. Moreover, such abatement activity would attain its maximum while the species population remained moderately high, and would *decrease* as the population declined and restoration neared. Even where abatement measures have been employed as a complement to restoration, they typically emerge only when the situation is already dire. Similarly, environmental policies towards hazardous waste tend to emphasize terminal cleanup and permanent storage (restoration) rather than slowing waste generation. The Harvard example is instructive here, as for years the university focused almost exclusively on hauling wastes away and largely overlooked methods for curbing their generation.

Regulation of air and water pollution, on the other hand, tends to focus on emissions rather than the resulting quality level in the environment. Where only flows matter, or restoration is unavailable, such an emphasis is optimal. But when pollution accumulates, policies should adjust if restoration is a possibility, either technologically or politically. For example, imagine that an economic technology is developed (some time in the future) to remove carbon dioxide from the atmosphere. In such a scenario, optimal abatement of carbon dioxide emissions would fall, and restoration would eventually take place if carbon dioxide levels climbed sufficiently.

In other realms, government may intervene in SFQ problems by directly affecting the balance between maintenance and restoration. For example, zoning regulations often treat new construction much differently than modification of existing structures, effectively imposing additional costs on restoration relative to maintenance. Similarly, environmental regulations that exempt existing sources of pollution from stringent regulations – a practice known as "grandfathering" – create incentives for firms to maintain old and inefficient plants rather than replacing them with new ones, distorting optimal decision-making.

While we have pointed out the parallels between our approach and the classic (S,s) model at points, we have not exploited them fully. A natural extension of this work, therefore, would be to introduce flow control into the standard (S,s) model. As we have already noted, that model can provide the basis for understanding capital investment and durable goods purchases; it also undergirds macroeconomic models of price adjustment (Sheshinski and Weiss 1977) and cash holdings of individuals (Baumol 1952; Tobin 1956). Incorporating flow control into the (S,s) model – following methods outlined here – could allow for a much richer theoretical description of a wide range of economic phenomena and public policy issues.

References

[1] Abel, A. B., and J. C. Eberly. 1994. "A Unified Model of Investment Under Uncertainty." *American Economic Review* 84: 1369-1384.

[2] ———. 1996. "Optimal Investment with Costly Reversibility." *Review of Economic Studies* 63: 581-593.

[3] Arrow, K. J., T. Harris, and J. Marschak. 1951. "Optimal Inventory Policy." *Econometrica* 19: 250-272.

[4] Baumol, William J. 1952. "The Transactions Demand for Cash: An Inventory Theoretic Approach." *Quarterly Journal of Economics* 66(4): 545-556.

[5] Boyle, Kevin J., and Richard C. Bishop. 1987. "Valuing Wildlife in Benefit-Cost Analyses: A Case Study Involving Endangered Species." *Water Resources Research* 23(5): 943-50.

[6] Caputo, M. R., and J. E. Wilen. 1995. "Optimal Cleanup of Hazardous Wastes." *International Economic Review* 36: 217-243.

[7] Dasgupta, P. S., and G. M. Heal. 1979. *Economic Theory and Exhaustible Resources.* Cambridge, UK: Cambridge University Press.

[8] Dixit, A. K., and R. S. Pindyck. 1994. *Investment Under Uncertainty.* Princeton: Princeton University Press.

[9] Falk, I., and R. Mendelsohn. 1993. "The Economics of Controlling Stock Pollutants: An Efficient Strategy for Greenhouse Gases." *Journal of Environmental Economics and Management* 25: 76-88.

[10] Feldstein, Martin S., and Michael Rothschild. 1974. "Towards an Economic Theory of Replacement Investment." *Econometrica* 42(3): 393-424.

[11] Grossman, S., and G. Laroque. 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica* 58: 25-51.

[12] Hotelling, H. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39: 137-175.

[13] Keeler, E., M. Spence, and R. Zeckhauser. 1971. "The Optimal Control of Pollution." *Journal of Economic Theory* 4: 19-34.

[14] Krylov, N. V. 1980. *Controlled Diffusion Processes*. New York: Springer–Verlag.

[15] Kushner, H. J., and P. G. Dupuis. 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. New York: Springer–Verlag.

[16] Nickell, S. 1975. "A Closer Look at Replacement Investment." *Journal of Economic Theory* 10: 54-88.

[17] Palmieri, A., F. Shah, and A. Dinar. 2001. "Economics of Reservoir Sedimentation and Sustainable Management of Dams." *Journal of Environmental Management* 61: 149-163.

[18] Phillips, C. V., and R. J. Zeckhauser. 1998. "Restoring Natural Resources With Destination-Driven Costs." *Journal of Environmental Economics and Management* 36: 225-42.

[19] Plourde, C. G. 1972. "A Model of Waste Accumulation and Disposal." *Canadian Journal of Economics* 5: 119-125.

[20] 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.

[21] 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.

[22] Smith, V. L. 1972. "Dynamics of Waste Accumulation: Disposal Versus Recycling." *Quarterly Journal of Economics* 86: 600-616.

[23] Stevens, Thomas H., Jaime Echeverria, Ronald J. Glass, Tim Hager, and Thomas A. More. 1991. "Measuring the Existence Value of Wildlife: What Do CVM Estimates Really Show?" *Land Economics* 67(4): 390-400.

[24] Tobin, James. 1956. "The Interest-Elasticity of Transactions Demand for Cash." *Review of Economics and Statistics* 38(3): 241-247.

[25] U. S. Fish and Wildlife Service. 1996. *California Condor Recovery Plan, Third Revision*. Portland, OR. 62pp.

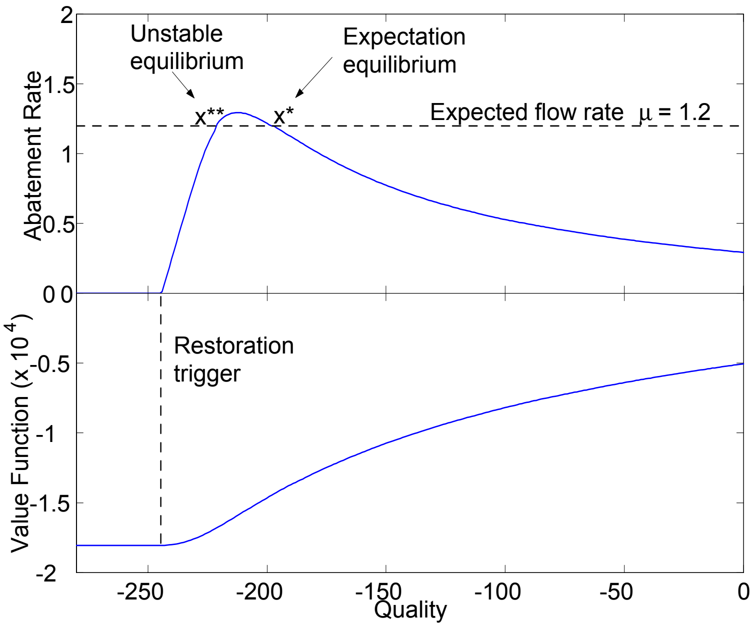

Figure 1: Optimal abatement path and corresponding value function. Note the different units of measurement on the positive and negative segments of the vertical axis.

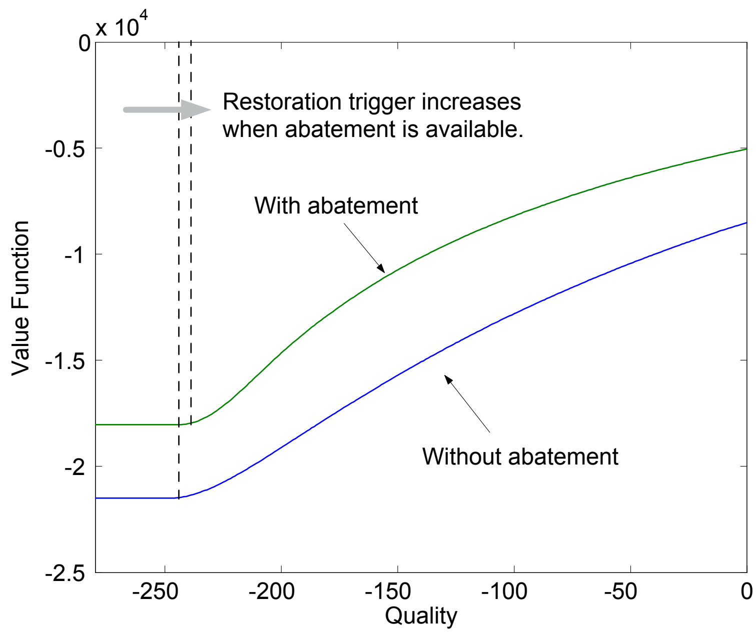

Figure 2: Value functions with and without the availability of abatement.

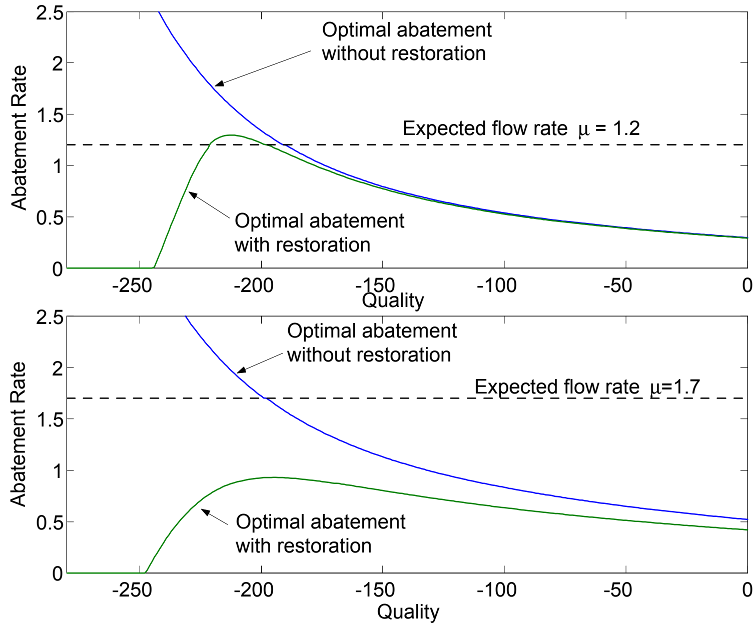

Figure 3: Effects of restoration on optimal abatement policy, for two flow rates.

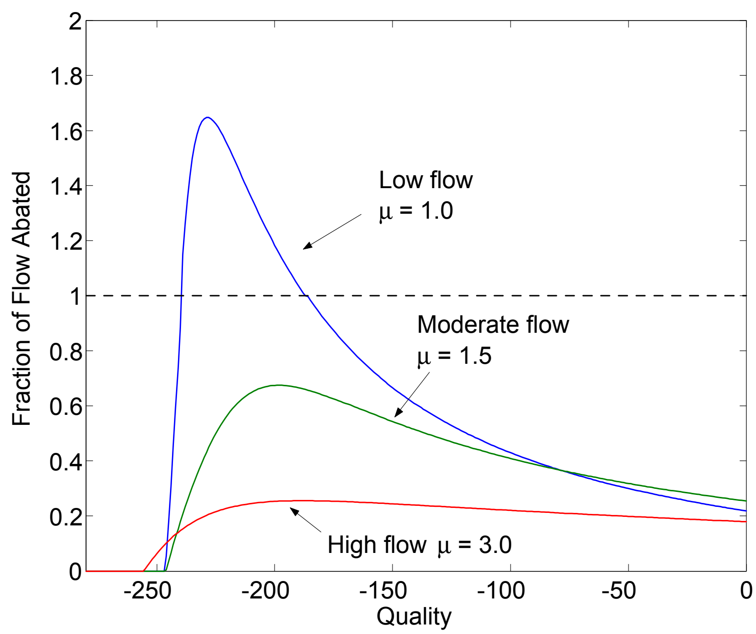

Figure 4: Fraction of flow optimally abated as a function of quality, for three flow rates.

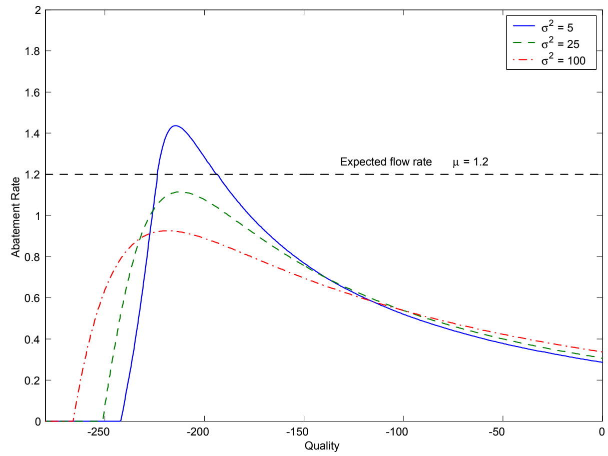

Figure 5: The effect of the variance rate σ^2 on the optimal abatement policy.

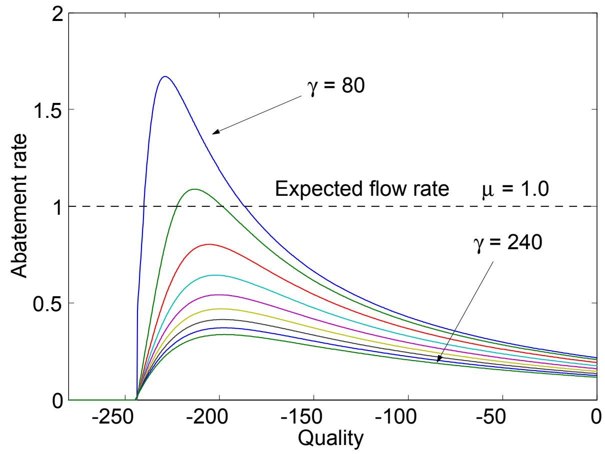

Figure 6: The effect of the marginal cost parameter γ on the optimal abatement policy..

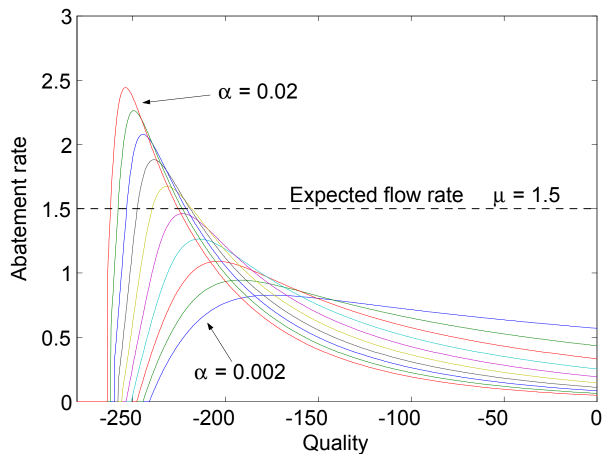

Figure 7: The effect of the discount rate α on the optimal abatement policy.

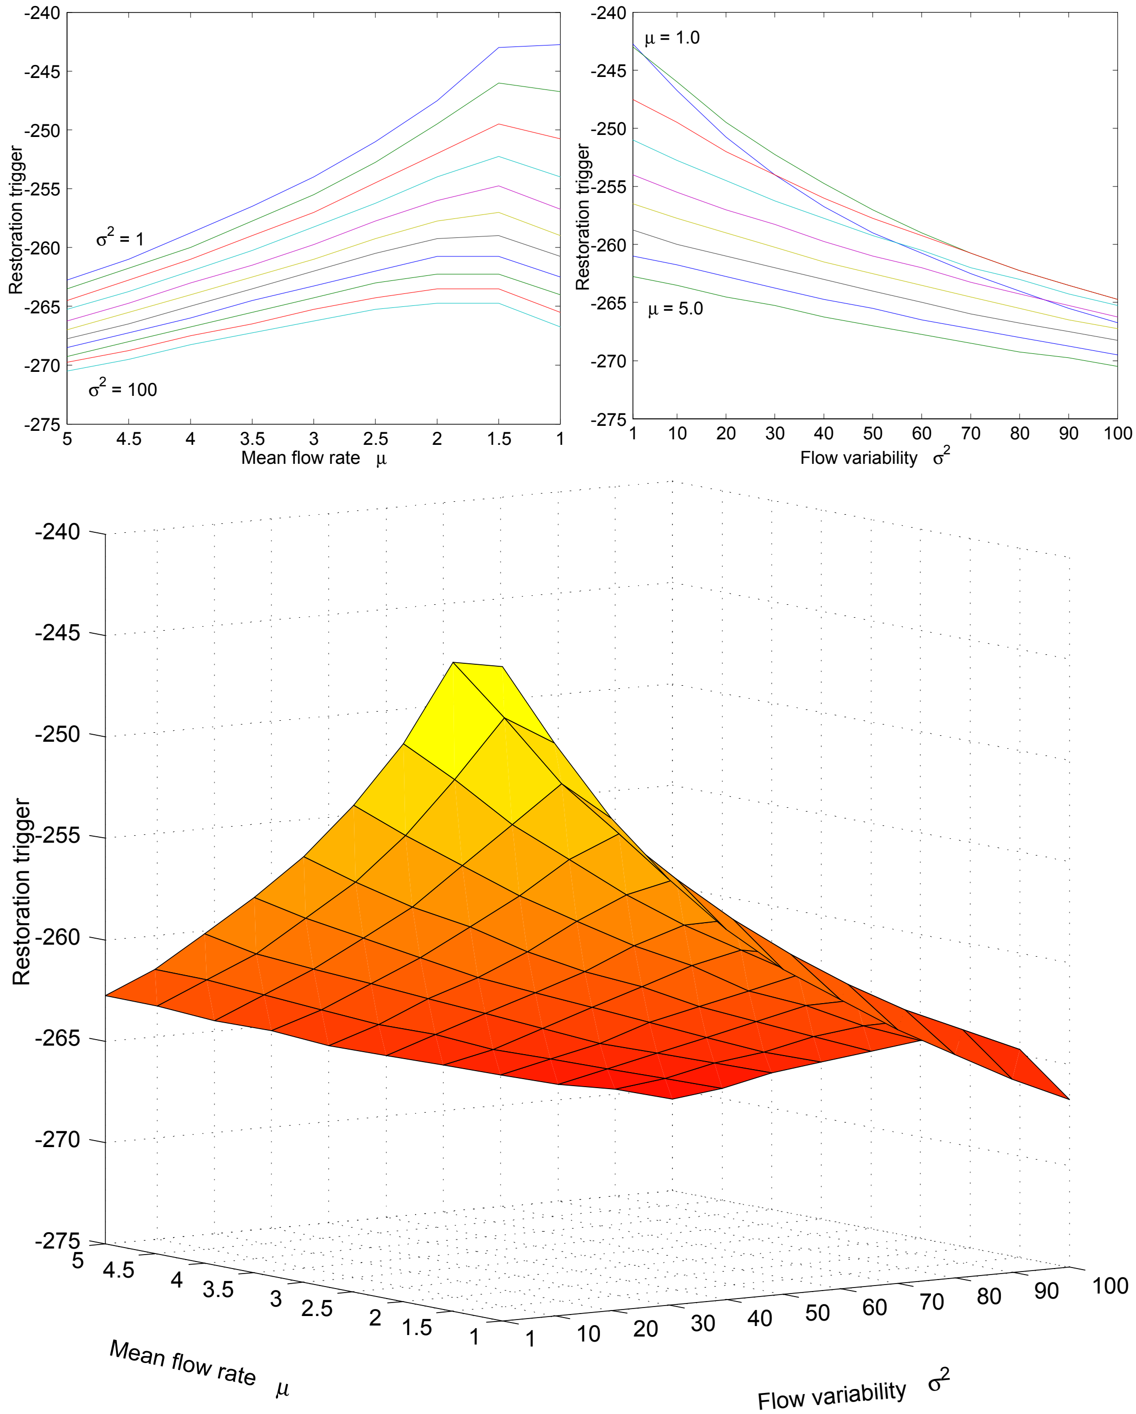

Figure 8: Optimal restoration trigger as a function of the mean flow rate and variability. In the bottom panel, both flow rate μ and variability σ^2 are depicted along the x and y axes, with μ ranging from 1 to 5 in increments of 0.5 and σ^2 ranging over (1, 10, 20, ..., 100). Hence the four sides of surface depicted in the figure run along the sides of the enclosing "box." The top two panels show the two-dimensional projections of the same surface.

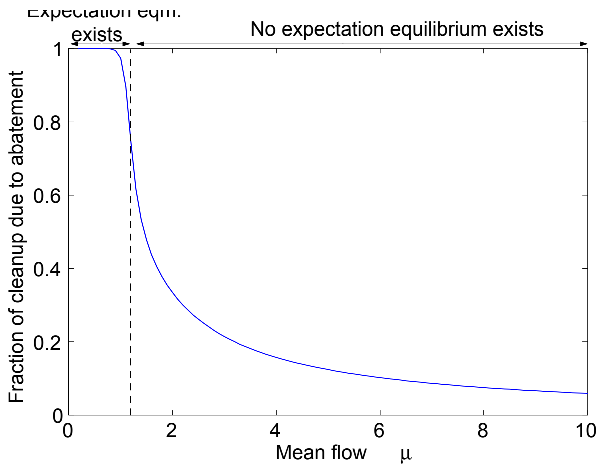

Figure 9: Fraction of optimal cleanup due to abatement, as a function of the mean flow rate μ.

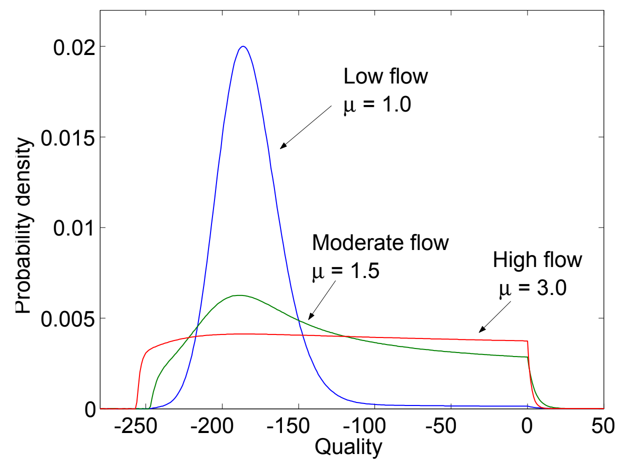

Figure 10: Frequency distributions of resource qualities (states) under optimal policies for three flow rates.

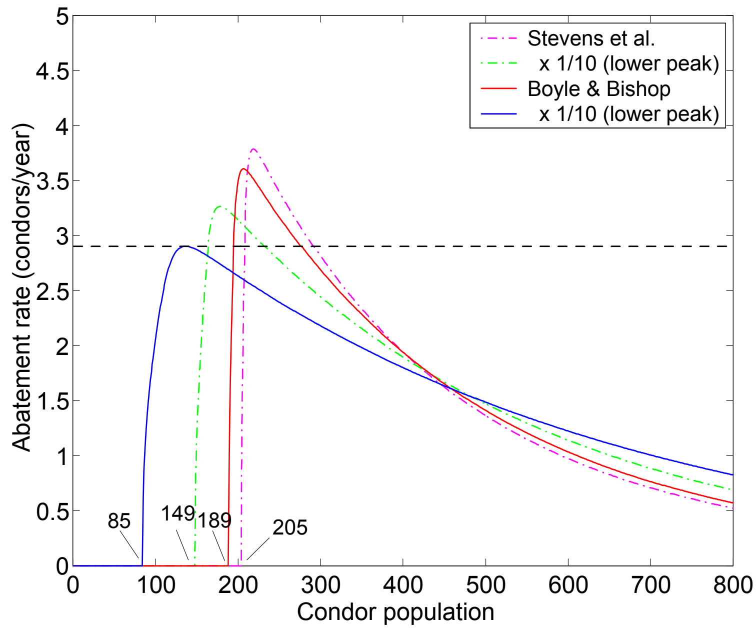

Figure 11: Simulated optimal abatement paths for the management of the California condor, for four estimates of willingness-to-pay. Restoration triggers are shown for each path.

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 $(\underline{x}, \infty)$ and differentiable everywhere. Furthermore, J satisfies

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

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 $(\underline{x}, \infty)$ (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

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

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 $a^*(x) = \overline{a}$, we have $a^*(y) > a^*(x)$. An entirely analogous argument establishes that $a^*(y) < a^*(x)$ if $x^\dagger < x < y$ and $a^*(y) \neq \overline{a}$. 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 $[0, \tau]$, where τ is given by

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

Let

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

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

Define "sample costs" associated with the three processes:

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

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

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

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

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

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

We consider two separate cases that together comprise a set of probability 1. The first is when

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

The second case is when $x_\tau^- = x^\dagger$. Given our assumptions on c , the fact that a^* is bounded above, and the fact that J is bounded and twice continuously differentiable on $(\underline{x}, \infty)$, it can be shown that for any $y > \underline{x}$, $|a^*(y) - a^*(y+\Delta)| = O(\Delta)$. It follows that $\sup_{t\in[0,\tau]} |a^*(x_t^-) - a^*(x_t^+)| = O(\Delta)$ and $x_\tau^+ - x^\dagger = O(\Delta)$. We then have

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

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

By Bellman's principle of optimality, we have

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

Hence,

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

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

q.e.d.

Proof of Theorem 2

Proof. For notational convenience, in this proof we shall write $\tilde{J}$ for J_{abate}, and likewise for $\tilde{a}$. 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 $[0, \overline{a}]$ 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

$$\sup_{z\in[0,\infty)} \tilde{f}_x(z). \tag{8}$$

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 $\tilde{a}(x) = \min(b(x), \overline{a})$. 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 $\tilde{a}(x) = \overline{a}$ for $x \leq \hat{x}$. Assertion (i) follows.

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

q.e.d.

Proof of Theorem 3

Proof. We continue to write $\tilde{J}$ for J_{abate}. 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

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

and

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

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

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

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

It follows that $J' < \tilde{J}'$, which gives us Assertion (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 $a^*(x) < \overline{a}$ then $\tilde{a}(x) > a^*(x)$. Hence, we have Assertion (ii).

q.e.d.

Proof of Theorem 6

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

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

$$= \lim_{T\to\infty} \frac{1}{T} E_x \left[\int_{t=0}^{T} u(x_t) dt \right] - C\mu\tilde{x},$$

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

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

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

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

is increasing in σ. Therefore,

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

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

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

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

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

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 $\overline{\sigma} > 0$, there exists some $\overline{\alpha} > 0$ such that for all $\alpha \in (0,\overline{\alpha})$,

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

and therefore

$$\frac{\partial J_{\text{restore}}(x,\sigma,\alpha)}{\partial \sigma}\Big|_{\sigma=\overline{\sigma}} > 0.$$

q.e.d.

Proof of Theorem 7

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

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

It follows that

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

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

Theorem 1 asserts that for any $C > 0$, $J(0, C) - J(\underline{x}(C), C) = C$. Inequality 9 therefore implies that $C_2 < C_1$, which yields a contradiction.

q.e.d.

Appendix B: Numerical simulations in Section 3

The computations that generated Figures 1-11 in Section 3 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 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, although more iterations were used in some cases.

Table B1. Parameter values and functional forms for figures.

variance rate		σ^2	=	9.0
discount rate		α	=	0.005
restoration cost		C	=	13000
abatement ceiling		$\overline{a}$	=	20
abatement cost	$c(a) = \gamma a^2$	γ	=	40
utility	$u(x) = -e^{-\beta x + \kappa}$	β	=	0.05
		κ	=	-7.5

Appendix C: Simulation of optimal condor restoration policy

Base simulations

Utility

Four estimates of per-capita annual willingness to pay were used, drawn from studies of willingness to pay for bald eagle populations, as discussed in the text: (i) Stevens et al. 1991 (\$32.94 for a population of 130); (ii) Boyle and Bishop 1987 (\$15.34 for a population of 300); (iii) Stevens et al. times 1/10 (\$3.29 for 130); and (iv) Boyle and Bishop times 1/10 (\$1.53 for 300). Each of these estimates was multiplied by 700,000 – the U. S. membership of the Sierra Club – to yield an estimate of total annual willingness to pay. We assumed that the willingness to pay for zero birds was zero. For utility functions throughout the condor simulations we used a negative exponential function of the form $-e^{-\beta x+\kappa}$, with the scaling parameter κ chosen to exhibit fairly rapid diminishing marginal returns. These assumptions then determined a value of β corresponding to each measure of willingness to pay. In the computations, the restoration point was normalized to be zero, and negative values of x were considered (as in the model of Section 3). For the purposes of exposition, we re-normalized the restoration point to 300 birds, and added the same number to every state in order to represent condor populations as positive numbers.

Cost

We estimated the costs of restoration and abatement using detailed program cost estimates in the Recovery Plan (U.S. Fish and Wildlife Service 1996). Recommended actions in the Plan were categorized as related to abatement or restoration. Restoration costs were simply summed (in present-value terms), resulting in an estimate of $C = \$15$ million to restore the population to a size of 300 birds (the target wild population that has been adopted by the Fish and Wildlife Service). For each abatement action, average cost (per condor per year) was estimated taking into account historical rates of decline in condor population and the priority accorded that action by the Service.[31] These unit costs were then arranged in increasing order, and a polynomial cost function was fitted to them. For the expected flow rate in the absence of abatement, μ, we used the average annual rate of decline of the condor population over the decades prior to restoration, based on wild condor censuses.

Restoration cost was determined from the 1996 Condor Recovery Plan (U.S. Fish and Wildlife Service 1996). The abatement cost function was generated by estimating unit costs for a number of protective actions recommended in the recovery plan and fitting a cubic polynomial to the resulting data. Estimates of the denominator ("abatement rate" or number of condors saved per year) were derived from the priority hierarchy attached to the various protective measures. Thus "Priority 1" actions, considered necessary to prevent extinction of the species, were taken to have greater effect than "Priority 2" actions (declared necessary to prevent significant decline) or "Priority

[31] Thus "Priority 1" actions, considered necessary to prevent extinction of the species, were taken to have greater effect than "Priority 2" actions (declared necessary to prevent significant decline) or "Priority 3" (necessary for full recovery). A table of the actions and associated costs is presented in Appendix C.

3" (necessary for full recovery). The abatement actions, along with their assigned priorities and associated unit costs, are presented in Table C1.

The specific functional forms and parameter values used in the simulation are summarized in Table C2. All dollar values are in constant 1996 terms. All costs are expressed in thousands of dollars.

Monte Carlo simulations

We performed 222 simulations where the utility parameter β, the cost parameter γ_2, and the variability σ^2 were the result of random draws from independent normal distributions. The random variables were generated by Matlab's pseudo-random number generator, with each draw representing three randomly chosen variables from ten supplied by the number generator. (In fact, we drew 225 parameter vectors and discarded three with negative values for the parameters; hence the distributions, strictly speaking, were truncated rather than truly normal. However, the truncation seems minimal enough that we can safely ignore its effects on inference from the simulations.) The parameters γ_2 and σ were drawn directly from the distributions given in Table C3. In the case of β the random parameter was per capita annual WTP, drawn from a normal distribution centered on \$7.70, equal to one-half the Boyle and Bishop estimate; this was chosen as a conservative estimate that still allowed ample room for variance. For each draw of WTP, say $\hat{u}$, β was computed by using the functional form $u(x) = -e^{\beta x+5}$, setting $u(300) - u(0) = \hat{u}$ (recall that 300 corresponds to the total bird population that the Boyle and Bishop study was based on), and solving for β.

We calculated the "average" abatement rate as the average over states for which abatement is positive. Note that it is computed over *states* themselves without taking into account the probabilities of those states being reached. We treated the average and maximum abatement and the trigger level $\underline{x}$ as three variables, and computed a joint 95% confidence region using the Wald statistic

$$W_i = (\boldsymbol{\theta}_i - \boldsymbol{\theta}_0)' \, \hat{\mathbf{V}}^{-1} \, (\boldsymbol{\theta}_i - \boldsymbol{\theta}_0) \, ,$$

where $\boldsymbol{\theta}_i$ is the 3×1 vector of the values of the three variables for draw i, $\boldsymbol{\theta}_0$ is the vector of sample means, and $\hat{\mathbf{V}}$ is the sample covariance matrix. The 95% confidence region was the set of draws for which $W_i < 7.82$, the critical value for the χ^2 distribution with three degrees of freedom.

The results of the Monte Carlo estimation are summarized in Table C4.

Table C1. Abatement measures and costs for condor simulation.

Priority[a]	Abatement measure	Annual cost ($000/yr)	Average cost ($000/condor-yr)
1	Protect habitat in S.W. Kern Co.	1	7
1	Protect habitat on Tejon Ranch	1	7
2	Protect habitat in Glenville Woody Area	1	8
2	Protect habitat near San Juan Creek	1	8
1	Provide information to private landowners	1	11
1	Protect nest sites	2	14
1	Protect roost sites	2	14
2	Manage condor foraging habitat	2	16
3	Protect habitat in Tulare County Grasslands	1	20
1	Provide information to land managers	2	21
1	Ongoing contaminant-related activities	20	30
2	Reestablish native ungulates	4	32
1	Protect Elkhorn Hill and Caliente Range areas	5	35
3	Provide dead livestock on rangelands	4	36
3	Assess historical findings	2	40
1	Distribute educational material	5	53
2	Establish observation points	5	53
3	Restrict aircraft	1	54
3	Perform federal land use planning	3	60
1	Protect habitat near Bitter Creek	10	70
3	Perform local and use planning	4	80
3	Hold training sessions	2	108
1	Protect habitat in Canizo and Elkhom Plains	20	140
1	Modify human-made structures (powerlines)	105	210
3	Step up enforcement of relevant laws	5	270
3	Protect habitat on Hopper Mtn.	20	400

Notes: a. Priority accorded abatement measures in the 1996 Recovery Plan from 1 (highest) to 3 (lowest).

Table C2. Parameter values and functional forms for condor simulations.

mean flow rate		$\mu = 2.9$
variance rate		$\sigma^2 = 9.0$
discount rate		$\alpha = 0.05$
restoration cost		$C = 15000$
abatement cost	$c(a) = \gamma_1 * a^2 + \gamma_2 * a^3$	$\gamma_1 = 0$
		$\gamma_2 = 4.223$
utility	$u(x) = -e^{-\beta x + \kappa}$	$\kappa = 5.0$
	Boyle & Bishop x 0.1	$\beta = 0.00704$
	Stevens et al. x 0.1	$\beta = 0.01018$
	Boyle & Bishop	$\beta = 0.01433$
	Stevens et al.	$\beta = 0.01720$

Table C3. Distributions for Monte Carlo simulations.

Parameter	Mean	Variance
γ_2	4.223	2.0
σ	3.0	2.0
WTP	$7.70	16

Table C4. Summary of Monte Carlo simulations: Outcomes in 95% confidence region.

Variable	Mean	Standard Deviation	Minimum	Maximum
maximum abatement rate	3.4	0.4	2.7	4.5
average abatement rate	1.1	0.1	0.9	1.4
restoration trigger	164	19	92	193