JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT **36,** 225–242 (1998)
ARTICLE NO. EE981051

Restoring Natural Resources with Destination-Driven Costs*

Carl V. Phillips[1]

University of Minnesota School of Public Health, Division of Environmental and Occupational Health, Box 807 Mayo, 420 Delaware Street, S.E., Minneapolis, Minnesota 55455

and

Richard J. Zeckhauser

Kennedy School of Government, Harvard University, Cambridge, Massachusetts 02138

Received January 2, 1997; revised August 6, 1998

Economics typically assumes that injured natural resources are restored along a fixed path of increasing marginal costs. By contrast, many restoration efforts—such as cleaning a contaminated aquifer or replacing the sand on an oil-tarnished beach—are characterized by *destination-driven costs*, which depend mainly on final quality, not the prerestoration quality. Given the resulting nonconvexities in cost, the optimal level of restoration may be a discontinuous and nonmonotonic function of post-injury quality. Regulatory rules should reflect these patterns, as should liability rules, since restoration plans and costs determine the expected cost of putting a resource at risk.

I. INTRODUCTION

Environmental regulation generally tries to balance the benefits of healthy natural resources, the benefits of actions that may injure them, and the costs of restoring them following injuries, as economic theory would prescribe. Such a theory, which is instructive for the design of effective policies [5], usually assumes continuity, diminishing marginal returns to resource quality, and increasing marginal costs of resource restoration. These assumptions yield a readily computed interior optimum for the level of effort that should be made when natural resources are injured. The optimum is typically displayed graphically as the intersection of marginal cost (of cleanup) and marginal benefit (from higher resource quality) curves.

This analysis departs from the traditional assumptions. It relies on two principal observations. First, in restoring many resources, the cost depends primarily on the level of restored quality, and little if at all on the amount of damage restored. We label this cost structure *destination-driven costs*. Second, destination-driven costs have particular implications for regulation. Indeed, policies based on conventional economic cost models will be inefficient when costs are destination-driven.

* The authors thank Miriam Avins, Henry Lee, John Lindsey, Albert Nichols, Rob Stavins, Paul Todgham, two referees, and an associate editor for helpful comments and Anne Marie Johnson for research assistance. Phillips thanks the Robert Wood Johnson Foundation for generous funding.

† Corresponding author.

0095-0696/98 $25.00

Increasing marginal costs of pollution reduction is often the appropriate model for reducing ongoing injuries from multiple sources, such as reducing a city's air pollution or the effluent in a large body of water. In such cases, the large number of possible sources of reduction, when considered in the efficient order, yield an approximately-continuous, upward-sloping cost curve for quality improvement.

The cost of remediating existing injuries at individual sites is far different, yet the increasing marginal cost model is often mistakenly extrapolated to that arena. For individual projects, such as Superfund sites, the remediation cost function is often poorly behaved and may offer increasing returns to scale for either restoration (cleaning up the injured resource) or replacement.

The engineering literature reports costs of cleaning up land and groundwater pollution in terms of the treatment technology and such things as the number of offset wells to drill, tons of earth to dig, or gallons of water to process. A remediation effort might consist of extracting water to evaporate volatile organic chemicals, digging up and washing the soil, introducing bacteria that break down pollutants, adding chemicals to chelate metals, or any of dozens of other techniques. In most cases, neither the cost nor the end result depends on the initial quantities of the pollutants, but only on the technology employed [8, 13, 20]. Methods to estimate the cost of a project discuss dozens of factors, but the concentration of contamination often plays little or no role (though it does have indirect effects like the level of protective gear that is needed and perhaps the size of the area polluted) [15]. Once remediation is underway, the same equipment must be used regardless of pollutant concentration, and the total cost and final outcome depend overwhelmingly on the choice of technology and its limits rather than the starting quality level. Such factors as the types of contaminants, soil composition, and moisture content affect the choice of technology. But the concentration of contaminants is not a major consideration in engineering discussions of technology choice, cost, or target outcome [7, 8]. In sum, engineering discussions contrast sharply with the usual formulation in the economics literature, making it evident that the cost of restoration is rarely a continuous, well-behaved function of either the total amount of injury or total amount of restoration.[1]

Current litigation (*U.S. v. Montrose Chemical Corporation of California et al.*) under the Comprehensive Environmental Response, Compensation and Liability Act (CERCLA) illustrates how the resource costs of remedies may be largely independent of the magnitude of injury. Actions by the responsible parties have contaminated sediments on the ocean floor in an area near Los Angeles with DDT and PCBs, causing reproductive problems to two species of fish, and to bald eagles and peregrine falcons. The government's current proposed remedy is that the responsible parties pay for capping (covering) the contaminated sediments, and for a breeding program for the birds.[2] The resource cost of these remedies would be little affected if there were half or twice as much contaminant in the affected area.

Replacing an injured resource, a second remediation approach, has similar implications. Replacement might involve removing contaminated topsoil and permanently storing it, replacing oil-covered sand on a recreational beach, or simply

[1] *Injury* refers to the physical alteration of the resource itself, in contrast to *harm*, which refers to the impact on human utility.

[2] The government is also seeking payment of hundreds of millions of dollars for lost use and nonuse values, a matter beyond the scope of this paper.

capping and abandoning a site and replacing it with another. If a water source becomes unsafe to drink, an alternative supply can be purchased; similarly, if a wetland is hopelessly injured from a toxic spill, another potential wetland can be recovered from agricultural use for some fixed cost [9]. Replacement will sometimes be more efficient than trying to salvage a particularly badly injured site, providing more net environmental gain for less cost [3]. When the dioxin contamination of Times Beach, Missouri was discovered in 1982, the immediate solution was for the Superfund to pay $33 million to relocate all residents and businesses to a safer location, a cost that was independent of the quantity of the pollutant [17, p. 386]. In such cases, the total cost of the effort will vary with some measures of the size of injury (total area polluted), but not with others (amount of contaminant per cubic foot). The final outcome is independent of the level of injury, but may have a substantial effect on the cost.

Consider an analogy outside the politically charged realm of natural resource injury, the case of a building that is injured by a fire. Up to a certain level of injury, the building will be salvaged and restored (though the last bit of smoke damage may not be repaired). This may follow the standard cost assumptions. But for injury past a certain level, the building will be torn down and replaced (possibly with a design that better meets modern needs).[3] A similar principle applies to automobiles —sufficient injury to a car will "total" it, resulting in its replacement. (The quality of the replacement may be higher or lower than that of the old car.) In these cases, as with optimal environmental restoration, the cost of remediation may be unrelated to the level of injuries over a wide range.

Popular discussions of environmental remediation fuel the economic assumptions that more concentrated pollution is more difficult to remedy, often describing pollution as being hundreds of times the maximum safe level or an area being unsafe to occupy for centuries. Despite the rhetoric, in most cases the cost of remediation is not significantly tied to the concentration of pollution. If regulatory and liability rules do not recognize actual patterns of costs, incentives will be wrong, and firms will make inefficient choices about which sites to put at risk and to what degree. Governments and firms responding to the malincentives will choose inefficient restoration efforts following injuries.

These lessons are particularly important for environmental regulations that deal with local remedies. Major relevant acts include the Superfund laws, CERCLA and the Superfund Amendments and Reauthorization Act (SARA), dealing with toxic waste sites; the Resource Conservation and Recovery Act (RCRA), dealing with contamination on actively-operated sites; and for some provisions (e.g., noncontamination in zones around aquifers) of the Clean Water Act.

This paper examines how restoration costs depart from standard economic assumptions, and demonstrates the efficiency implications. Section II formalizes our description of restoration costs, which allows for destination-driven projects, and contrasts it with the standard assumptions. Section III discusses the optimal restoration decision under the standard formulation, presents the contrasting results given our alternative cost formulation, and compares the two. Section IV discusses implications for decision making and regulation. Section V concludes.

[3] Public policy decisions about improving infrastructure encounter similar nonconvexities. For example, a municipality might want to expand its airport to reduce limited overcrowding, but congestion at some critical level may call for a whole new facility.

II. COST FORMULATIONS FOR THE RESTORATION DECISION

Let $V(q)$ be the total discounted value of the stream of benefits a resource produces at quality level q.[4] Define s as the starting quality of the resource. Define t as some higher target quality level to which the resource could be raised (and assume for concreteness that this is restoration following some injury). Let $C(s,t)$ be the dollar cost of restoring the resource from quality s to quality t.

The Standard Assumptions

It is usually assumed that

$$\frac{\partial V(q)}{\partial q} > 0, \tag{1}$$

$$\frac{\partial^2 V(q)}{\partial q^2} < 0 \tag{2}$$

and

$$\frac{\partial C(s,t)}{\partial t} > 0, \tag{3}$$

$$\frac{\partial^2 C(s,t)}{\partial t^2} > 0 \qquad \text{for } s \text{ constant}. \tag{4}$$

The first pair of inequalities, (1) and (2), reflects the assumption that the value of a resource increases with quality but that the marginal value is diminishing, and imply that the loss resulting from an injury to a natural resource is convex in the level of injury. The second pair of inequalities, (3) and (4), reflects the assumption that the marginal cost of restoration is positive and increasing in the target quality level (holding the pre-restoration level constant). These assumptions result in a downward-sloping benefits curve and, for the given s, an upward-sloping cost curve. They ensure that the optimal level of restoration for a resource will be unique (possibly zero), and that it is seldom optimal to restore to the maximum possible quality.

In assuming these conditions explicitly, it is also usually implicitly assumed that there is some function of quality, $c(q)$, such that

$$C(s,t) = \int_s^t c(u)\, du, \tag{5}$$

with $c(q)$ positive and increasing. That is, it is assumed that there is some marginal cost function that describes the cost of incremental restoration at any given quality level, independent of s or t, that meets conditions (3) and (4). Drawing a marginal

[4] Utility can change as a result of how a situation is framed [18] and the change in the quality of a resource, rather than just the absolute level, might matter [19]. Determining the shape of the curve (that is, measuring the value of natural resources) is also problematic [11]. However, since the present analysis is normative in nature, it seems reasonable to assume the existence and measurability of V, despite the widely documented anomalies and measurement difficulties.

cost curve and using it to discuss different possible levels for s and t implicitly assumes there is a $c(q)$ as defined by (5). We call the condition that there is such a $c(q)$, and in particular that it is independent of s, *initial-quality-invariant marginal costs* or *IMC*.

It follows from the existence of $c(q)$ as described in (5) that

$$\frac{\partial C(s,t)}{\partial s} < 0 \qquad \text{for } t \text{ constant}. \tag{6}$$

Both (3) and (6) can be interpreted as "the bigger the restoration, the more it costs," but the two conditions are fundamentally different. Condition (3) says that the higher the target level of restoration, the higher the cost (for the same starting level); condition (6) says the higher the starting level, the lower the cost (to get to the same target level). This distinction is important since the present analysis discusses important and common cases that are not IMC and for which (6) is false but (3) is true.[5] In such cases, the desirable policies deviate substantially from conventional prescriptions, and would encourage greater risk in some cases and greater caution in others.

Departures from the Standard Assumptions

Several classes of departures from the standard convexity assumptions in terms of benefits from environmental quality and the accumulation of pollutants are important. Leaving restoration possibilities aside, marginal harm (measured in human or ecological suffering) from an injury to a natural resource may not increase with the magnitude of injury. For uninjured resources, the first unit of pollution may have a much higher marginal cost than subsequent units. For example, a light coating of oil on a beach endangers the native fauna and flora and dramatically diminishes the recreational and aesthetic value; doubling the volume of oil would far less than double the impact. Similarly, ten pieces of litter on a mountain trail diminish the wilderness experience substantially more than one-third as much as thirty pieces would [12, p. 131].

When damages are severe, marginal losses also may not be convex. For major injuries, there is likely to be a point beyond which the resource has very little human or ecological value left to lose [6]. At that point, further quality reduction would entail little further loss, implying diminishing marginal costs. Moreover, people can engage in averting behavior and avoid a resource with lowered quality. Human activity will move elsewhere when quality drops below a certain point, even if some value remains [14]. For example, once there is enough litter on a trail, people will choose another trail, and no one will swim at an oil-coated beach, even if the coating is light. For people who are no longer using a resource, additional injury causes no loss of benefit. The reduction of pollution due to natural processes may be a nonmonotonic function of the total pollution stock, contrary to the simplified formulation that assumes that the reduction is convex and increasing in

[5] Strictly speaking, in the cases we are discussing, condition (3) may not hold, but a very similar stepwise inequality will. (This is discussed further in footnote 7.) Also, the present discussion applies to cases where (6) is true but $c(q)$ approaches zero as s decreases. Details of these points are omitted for clarity of exposition.

pollution stock [16]. This may make the marginal costs of pollution nonmonotonic over the level of pollution.

The present analysis focuses on nonconvexities on the restoration side of resource injury. For restoration of individual sites, the IMC condition is frequently not met (perhaps almost never), resulting in economies of scale and other violations of the standard model. Economists tend to think of restoration projects as being incremental. Total cost is the sum of incremental cost steps, and the project can stop at any step. This view may result from an analogy to multiple-site or multiple-source environmental injuries, where substantial injury reduction comes as a result of many small projects which are assumed to be undertaken in the order of diminishing efficiency. Under those circumstances, conditions (3)–(5) hold and all restoration projects follow the path traced by the same marginal cost curve. Local optimization efforts lead to the global optimum.

But many remediation efforts consist of a single restoration project, carried out at the only efficient size, and achieving a particular t, regardless of s. The cost depends almost entirely on terminal (post-restoration) quality, contrary to condition (6). Thus, costs are destination-driven; they are characterized by *invariant total costs*, or *ITC*. When an injury is remedied by replacement, whether the replacement is a single water well or an entire town like Times Beach, ITC will also apply.

The significance of injury magnitude depends on how the injury will be remedied. If an ITC project will be employed, then further injury to a site would neither change the optimal choice of restoration project nor add to its cost, which has significant implications for policies to control injuries. The potential use of ITC projects means that the optimal response to an injury may be discontinuous in the level of the injury. A small increase in injury, resulting in a lower s, may result in a discontinuous increase in the optimal post-restoration quality and expenditure, as we show below.

Naturally, the cost of most projects in the real world will not be described perfectly by either IMC or ITC. The potential options for restoring sites may have characteristics of both. This would occur for a largely destination-driven project, for example, if additional injury creates some additional cost of removing a greater volume of material or requires greater safety precautions for workers during cleanup. It would also occur if the cost of stabilizing the injury (e.g., to prevent geographic expansion) increased even if the cost of the eventual restoration did not. For simplicity, this analysis focuses on the polar cases of pure IMC [where (3)–(6) all hold] and pure ITC [where (5) and (6) are false].[6]

[6] Many of the qualitative results hold under weaker conditions, and almost-ITC projects remain different from any almost-IMC formulation. The stronger conditions are consistent with a description of the real-world situation and allow for clearer exposition. For example, if $\partial C(s,t)/\partial s$ only approaches zero but remains negative, the qualitative results emphasized in the analysis would remain valid, though the quantitative details would be more complicated. The critical points and discontinuities would still exist, but their positions would be more difficult to calculate and the results more difficult to demonstrate. Similarly, it may be more realistic to model the cost of restoration as a step function anchored at s, rather than continuous incremental costs. Its implications remain close to those for IMC, and are very different than those for ITC.

III. OPTIMAL RESTORATION DECISIONS

IMC Projects

Projects with IMC have the standard continuous incremental properties that are traditional in economics. It follows immediately from (5) that

$$C_I(s,t) = C_I(s,r) + C_I(r,t) \qquad \text{for all } s \le r \le t, \tag{7}$$

where $C_I(s,t)$ denotes the cost of going from s to t using only incremental projects. There is effectively no difference between a single incremental restoration from s to t and the combination of two projects that first restore from s to r and then from r to t.

If a resource is injured, dropping its quality level to s, we wish to find the socially efficient target restoration level, t, given the cost function from Eq. (5). Let $B(s,t)$ be the benefit from a restoration effort that improves the quality of a resource from a starting level s to a target level t. Then $B(s,t) = V(t) - V(s)$.[7] We accept assumption (1), that $V(q)$ is strictly increasing in q, and thus $B(s,t) > 0$ for all $s < t$. We wish to find

$$\underset{t}{\text{Max}}\,(\text{Net Benefits}) \equiv \underset{t}{\text{Max}}\,(B(s,t) - C(s,t)). \tag{8}$$

If assumption (2) about the diminishing marginal value of quality applies, the familiar process of equating marginal costs and benefits yields the unvarying interior optimum (assuming that escalating marginal costs prevent a pristine outcome from being optimal). Then (8) has a unique solution, $t = q_e$, defined by $\partial V(q_e)/\partial q_e = \partial C(s,q_e)/\partial q_e$. If an injury brings the resource below quality q_e, it should be restored to that level. Above q_e, it is efficient to undertake no restoration. Thus, if the IMC condition applies, there are two possible efficient final states: the initial quality level, s (with no restoration undertaken), and q_e.

ITC Projects

Label the set of all destination-driven projects **D**. For these projects, under pure ITC,

$$C_D(s_1,t) = C_D(s_2,t) \qquad \text{for all } s_1, s_2 \le t, \tag{9}$$

where $C_D(s,t) \equiv C(T) > 0$ is the cost of T, the least costly element of **D** that achieves state t (which is defined if and only if there exists an element of **D** that

[7] Note that $B(s,t)$, like $C(s,t)$, is measured in terms of total present value. In particular, since V is the discounted sum of the stream of all future values, B is the sum of added utility over all future periods.

This analysis does not separately address the stream of lost benefits that occurs between the time of an injury and the completion of a restoration project. While such losses may be substantial, some portion of the costs exist for any restoration project (including doing nothing), and thus fall out of the optimization decision. To the extent that interim costs vary across restoration projects, the difference can be considered part of the cost of the project.

achieves state t).[8] This implies that for all such t there is a function $C_D(t)$ such that $C_D(s,t) \equiv C_D(t)$ for all $s < t$. It follows immediately that, in contrast with (7),

$$C_D(s,t) < C(s,r) + C_D(r,t) \qquad \text{for all } s < r < t, \tag{10}$$

implying that there are economies of scale over some range.

In general, there may be multiple elements of **D**, including discrete options and continuous ranges. For example, to restore a wetland that is polluted with toxins, we could remove the contaminated water and soil and try to recreate a working wetland; we could remove the contaminants but fail to rebuild the wetland; or we could just pave the mess over and build a new wetland elsewhere (with a range of possible expenditures at the new site). We can simplify our analysis of ITC projects, however, by observing that there will be an element of **D** that is weakly superior to other elements of **D**, given the functions V and C_D. Formally this yields the following proposition.

PROPOSITION 1. *Let* $\mathbf{D} \neq \emptyset$ *and* $\mathbf{Q}$ *be the set of quality levels attained by some element of* $\mathbf{D}$. *Assume* $\mathbf{D}$ *is such that there exists* $\mathbf{E} \subseteq \mathbf{D}$ *consisting of projects that generate every* (t, c), *where* $t \in \mathbf{Q}$ *and* $c = \inf\{C(D)\colon D \in \mathbf{D}$ *and D results in quality* $t\}$, *and that set* $\mathbf{E}$ *is closed with respect to t and c. Then for any s, there will be a project, P, in* $\mathbf{D}$ *that produces target quality* $q_P \geq q_e$ *at cost* $C(P)$, *such that the net benefit* $B(s, q_P) - C(P)$ *is weakly greater than the net benefit for any other element of* $\mathbf{D}$.

A proof is given in the Appendix.[9] Note that **E** is simply the set of all target quality levels and the cheapest project that attains them. The practical implication of Proposition 1 is spelled out in Corollary 1.

COROLLARY 1. *In searching for the most efficient restoration option, we can treat the menu of destination-driven options as a single option, P, that yields quality* $q_P \geq q_e$ *at a cost of* $C(P)$.

Corollary 1 will be used implicitly throughout the analysis that follows, by replacing the menu of destination-driven options in the optimization decision with

[8] $C_D(s,t)$ may not be defined for all t since there may not be a destination-driven project that gets to a given t. This relates to the previous discussion that inequality (3) does not hold if there are destination-driven costs, but it is clear that an inequality with similar implications remains true. Specifically, $\partial C(s,t')/\partial t \geq 0$, where $t' = \operatorname{argmin}_T\{C(s,T)$ such that $T \geq t\}$. If only destination-driven projects are possible, this statement says that as t increases, $C_D(s,t')$ for the least expensive $t' \geq t$ for which $C_D(s,t')$ is defined increases weakly. If both destination-driven and incremental projects are possible, then the cost of getting to the cheaper of t or the least expensive $t' \geq t$ for which $C_D(s,t')$ is defined increases weakly with t.

[9] The second sentence of the proposition simply says that D is closed in relevant directions, and in particular contains a minimum cost project for each possible target quality. This assumption is of little practical import. If there were a series of possible destination-driven projects that approached as a limit some quality or cost, but never reached it, it would be a mathematical curiosity, but from a practical standpoint we could choose something close and call it the limit.

The proposition actually holds under slightly weaker, though less intuitive, conditions. Specifically, it is sufficient that $\mathbf{D} \neq \emptyset$; the set $\mathbf{Q}$ is bounded above; $C(D) \geq 0$ for all $D \in \mathbf{D}$; and for any sequence of D_i, $i = 1, 2, \ldots$ with $\lim_{i \to \infty} q_{Di} = q$ and $\lim_{i \to \infty} C(D_i) = c$, there exists D with $q_D \geq q$ and $C(D) \leq c$.

the single project P. If we consider only destination-driven projects, then there will be two candidates for the optimal post-restoration quality: s (which will result if no restoration is undertaken) and q_P.[10]

Comparing IMC and ITC Options

If both incremental projects and destination-driven projects are possible, then after an injury that brings quality to s, there are three possible efficient post-restoration quality levels: s, q_e, and q_P. Figure 1 illustrates the optimization decision. The curves $C(s_1, q)$ and $B(s_1, q)$ respectively denote the total costs and benefits of incremental restoration projects as a function of the final quality, starting from initial quality $s = s_1$.[11] If destination-driven projects are not available, then for any $s < q_e$ incremental restoration to q_e should be carried out. At target level q_e, the vertical distance between curve C and curve B is greatest (that is, net benefits from incremental restoration are maximized). This distance, the maximum net benefit from an incremental project, is labeled NB_I. The optimal destination-driven project, P, yields net benefits $B(s_1, q_P) = NB_D$. If both types of projects are available, these respective maximum net benefit levels must be compared. In the case illustrated in Fig. 1, this comparison favors the destination-driven project.

[10] Much of this analysis relies on the assumption that further injury to the resource is improbable. If there is a high enough probability of further injury, then it might be efficient to delay restoration under either IMC or ITC. With ITC, the optimal destination-driven project may result in quality higher than q_P, to compensate for expected future quality loss.

[11] The curves shown are for linear marginal costs and benefits (and thus quadratic total costs and benefits), as in a typical textbook model. The results apply to any cost and benefit curves that meet conditions (1) through (6).

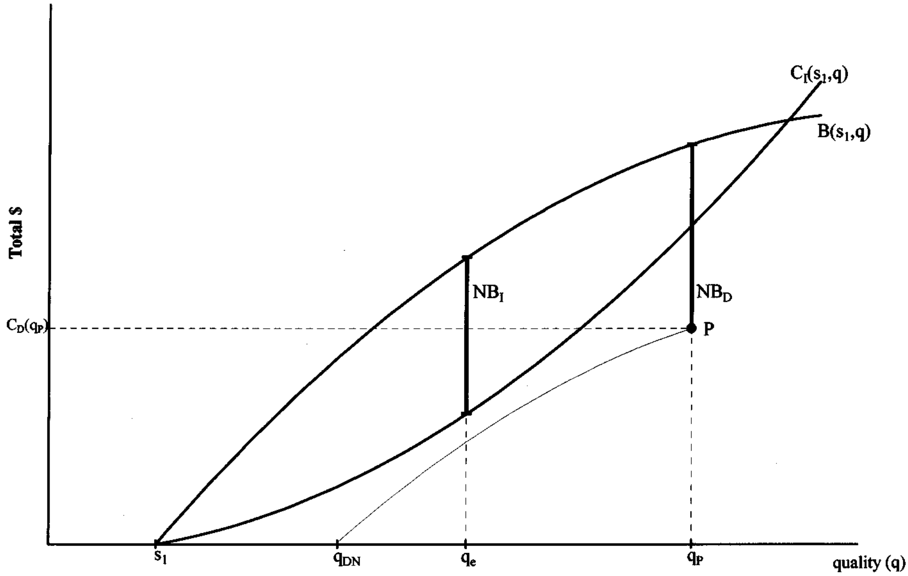

FIG. 1. Comparison of incremental and destination driven projects.

The general solution for the optimal choice of restoration project will detail the ranges of quality over which the various choices are most attractive. Incremental projects are more attractive than doing nothing if and only if $s < q_e$. The optimal destination-driven project will be more attractive than doing nothing if and only if $B(s, q_P) > C(P)$. Since $\partial B(s, q_P)/\partial s < 0$ and $\partial^2 B(s, q_P)/\partial s^2 \leq 0$, there will be a threshold, q_{DN}, such that for $s < q_{DN}$ project P yields positive net benefit, implying destination-driven restoration is superior to doing nothing. Specifically, q_{DN} is defined such that $B(q_{DN}, q_P) = C(P)$, or equivalently, such that the point $(q_P, C(P))$ lies on the total benefit curve for restoration originating at q_{DN} (as illustrated in Fig. 1). (The point can be identified in closed form as $V^{-1}[V(q_P) - C(P)]$, where V^{-1} is the inverse of the value function, $V(q)$, which must exist since $V(q)$ is strictly monotonic.)

Project P is more attractive than incremental restoration if and only if

$$[B(s, q_P) - C(P)] - [B(s, q_e) - C_I(s, q_e)] > 0, \tag{11}$$

which is equivalent to

$$C_I(s, q_e) > C(P) - B(q_e, q_P), \tag{12}$$

since $B(s, q_P) - B(s, q_e) = B(q_e, q_P)$ by the definition of B and $q_P \geq q_e$.

The right-hand side of (12) is constant over s, and the left-hand side is decreasing in s.[12] If the inequality is satisfied for any s, it will hold for all s less than some q_{DI}, the threshold for preferring the destination-driven project to the incremental project, defined such that $C_I(q_{DI}, q_e) = C(P) - B(q_e, q_P)$, if such a q_{DI} exists. (If incremental costs fall low enough at low quality levels, the destination-driven project may never be attractive.)

What are the possible orderings of q_e, q_{DN}, and q_{DI}, and how we can interpret them? There are two possible cases, since the ordering of q_e and q_{DN} determines as well the ordering of q_{DI}. The first case is $q_{DN} < q_e$. Given that $q_{DN} < q_e$, it must be that $q_{DI} < q_{DN}$, since for $s = q_{DN} - \epsilon$, the benefit from P is infinitesimal while the benefit from incremental restoration to q_e is first order. For this case, we would want to do project P if and only if there were an injury to the resource sufficiently large to drive quality below q_{DI}. Otherwise, ongoing incremental restoration would be efficient, and q_{DI} would never be reached. A second possibility is $q_e < q_{DN}$ (in which case q_{DI}, as defined above, does not exist). In this case, project P is always more attractive than incremental restoration. The site should be restored to q_P if and only if $s < q_{DN}$; otherwise it should not be restored.

To illustrate this, consider Fig. 1, which is characterized by $q_{DN} < q_e$ (and thus $q_{DI} < q_{DN}$). We can see that this inequality holds by observing the benefit curve that runs through point P crosses the horizontal axis above s_1. As observed earlier, at s_1 project P is the most attractive option, which implies that $s_1 < q_{DI}$, so q_{DI} lies between s_1 and q_{DN}. Figure 2 denotes the same functions with an alternative starting point, s_2. The figure demonstrates the graphical derivation of q_{DI}, given by finding the C_I curve that passes through the point $[q_e, C(P) - B(q_e, q_P)]$. The ranges delineated below the figure show the optimal project as a function of the pre-restoration quality level. From the position of q_{DI}, it is clear that the incre-

[12] The right-hand side of (12) is identically equal to the constant $B(q_{DN}, q_e)$, since $B(q_{DN}, q_P) = C(P)$, $C(P) - B(q_e, q_P) = B(q_{DN}, q_P) - B(q_e, q_P) = B(q_{DN}, q_e)$.

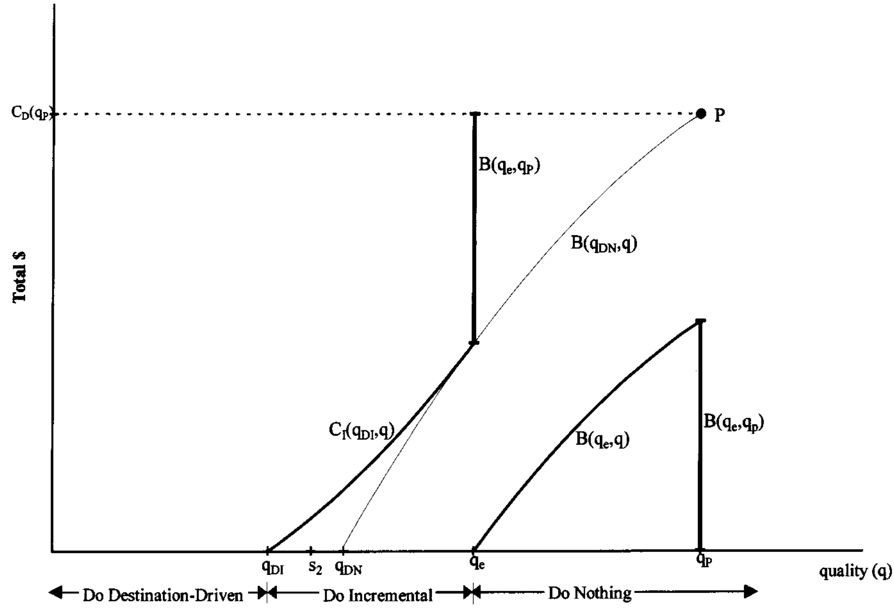

FIG. 2. Ranges of optimal restoration projects for $q_e > q_{DN}$.

mental project is preferable for initial quality s_2, even though the destination-driven project still has positive value there.

For the parameter values in Figs. 1 and 2, Fig. 3 illustrates how the optimal t changes as a function of s. For starting points above q_e, the site should not be restored, and so $t = s$. Between q_{DI} and q_e, incremental projects produce a positive net benefit, and this benefit is greater than that from project P. Thus the resulting quality level is q_e. Below q_{DI}, project P is optimal, and the resulting quality level will be q_P. This makes the optimal level of t nonmonotonic in s, and creates a discontinuity with the optimal target level jumping up once the post-injury level drops below q_{DI}. It follows immediately from this that policies that set a single standard for restored quality level would be inefficient even if the standard were set based on perfect knowledge of restoration costs and benefits.

To illustrate the other possible ordering, consider Fig. 4, which shows a different hypothetical destination-driven project. The new, inexpensive project P is a much better deal than any incremental project; hence $q_{DN} > q_e$. For initial state s_3 or any other point below q_{DN}, project P is optimal. Above q_{DN}, doing nothing is optimal.

Proposition 2 distills the discussion above.

PROPOSITION 2. *If both incremental and destination-driven projects are possible, then*: (1) *For* $q_{DN} < q_e$, *doing nothing will be optimal for* $s > q_e$; *incremental restoration to* q_e *will be optimal for* $q_{DI} < s < q_e$; *and destination-driven restoration to* q_P *will be optimal for* $s < q_{DI}$ (*if* q_{DI} *exists*). (2) *For* $q_{DN} > q_e$, *doing nothing will be optimal for* $s > q_{DN}$, *and destination-driven restoration to* q_P *will be optimal otherwise*.

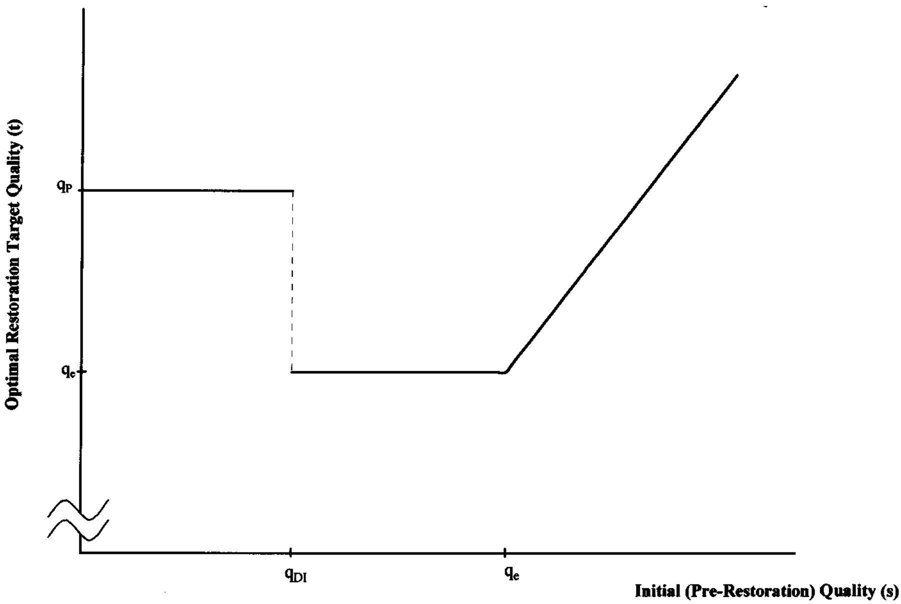

FIG. 3. Optimal target quality as a function of prerestoration quality for $q_e > q_{DN}$.

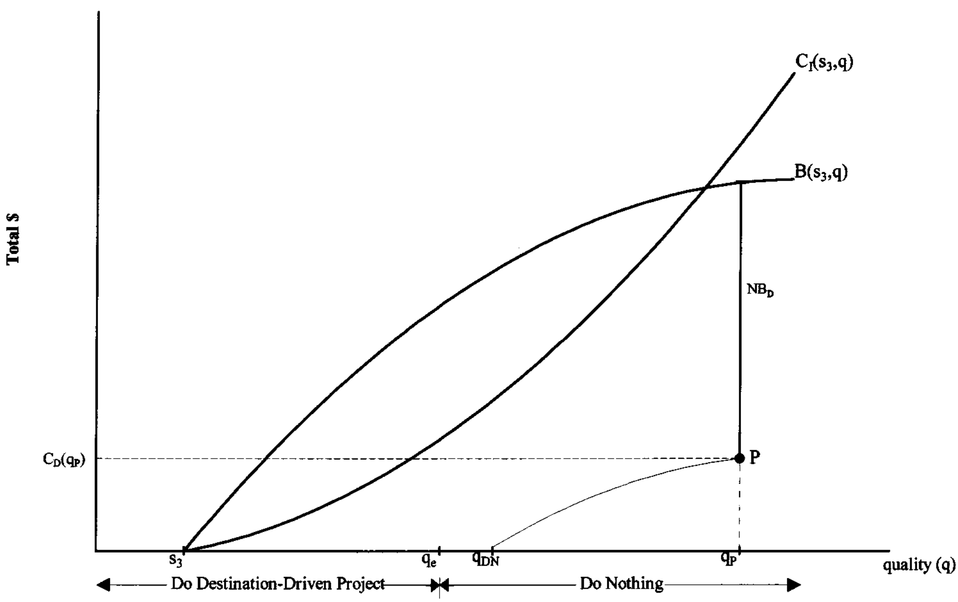

FIG. 4. Ranges of optimal restoration projects when $q_{DN} > q_e$.

Corollary 2, which follows immediately from Proposition 2 and the preceding discussion, shows that in a world with ITC projects there are several major departures from the standard efficiency analysis based on IMC projects.

COROLLARY 2. *If both incremental and destination-driven projects are possible, then the standard results, which assume convex costs of restoration, will not hold. In particular*:

(i) *If it is optimal to undertake some restoration project, the optimal target level, t, may vary with s (in particular, sometimes it will be q_e and sometimes it will be q_P).*

(ii) *Following a decrease in quality from pre-injury level q_0, where no restoration was efficient, the efficient level of t may be greater than q_0.*

(iii) *The marginal cost of additional injury may be zero once quality has dropped to a certain level because the desirable restoration does not increase in cost.*

The proofs of the proposition and corollary are implicit in the preceding discussion. Figure 3 illustrates parts (i) and (ii). These results are contrasted with the IMC model in Table I. Some further details of the proofs appear in the Appendix.

TABLE I

Properties of IMC and ITC Restoration

	IMC	ITC
Total cost is . . .	. . . the sum of marginal costs over the restoration range; highly dependent on starting quality	...fixed for a remediation or replacement option with a given target quality; independent of starting level
Most applicable scale	Many sources contributing to a widespread, ongoing injury	Local, self-contained, existing injuries
Example of relevant injury	Airborne sulphur emissions	Soil and groundwater pollution from an industrial facility or spill
Example of relevant legislation	Clean Air Act	CERCLA, SARA, RCRA
Optimal restoration level	A fixed q_e when damage brings quality below q_e, otherwise no restoration; monotonic in s	Depending on post-injury quality level, q_P may be preferable to q_e; may be nonmonotonic in s; may be higher than pre-injury level
Implicit marginal cost of further injury	Always positive and diminishing with current quality	May be zero
Implication for concentrating injury	Some incentive, but limited	Greater incentive
Net quality change following injury and efficient restoration	Negative or zero	Negative, zero, or positive
Effect of net quality change on damage awards	RP must pay residual loss of value	RP may receive credit for net increase in value

Part (iii) of the corollary points out that it may be desirable to allow multiple injuries to accumulate at a site, even if immediate restoration would produce net benefits.[13] A higher probability of future injury increases the incentive to wait before restoring. It may be efficient to allow the probability to increase, such as by locating a polluting industry at an injured site rather than restoring the site while locating the industry elsewhere. Multiple injuries will cost little more to restore than one, while restoring each as it occurred would be considerably more expensive. This savings trades off against the cost of living with lower than optimal environmental quality until the cleanup is undertaken.

IV. IMPLICATIONS FOR RESTORATION DECISIONS, REGULATION, AND LIABILITY

Restoration Decisions

When the costs of cleaning up pollution or restoring natural resources are primarily destination-driven, the desirable level of restoration depends on the current state of a resource. After an injury, it may be desirable to do nothing, do partial or full restoration, wait for further injury even though there is already a net beneficial project available, or undertake remedies that improve the resource's condition beyond its initial level. When the last option is chosen, a net improvement results from the injury and restoration. For example, when a vacant city lot gets sufficiently strewn with tires, it may attract attention and be turned into neighborhood gardens. A beach that has long suffered minor pollution from offshore bilge cleanings may become cleaner if a tanker spill induces a major cleansing or replacement.

By contrast, most current regulatory regimes seek to restore injured resources to exogenously defined levels of cleanliness or quality, generally at or below the pre-injury quality. Such an approach is warranted (assuming the right target level was specified for each particular resource) if cleanup costs are IMC. If there had been a prior incremental restoration, the pre-injury quality level would now be optimal as a target. However, always restoring to a single standard or always restoring to preinjury levels is a poorly chosen policy if cleanup costs are sometimes ITC or otherwise exhibit substantial economies of scale, as suggested by parts (i) and (ii) of Corollary 2.

Regulatory Harm Prevention

Various regulatory systems limit the injuries to natural resources from otherwise economically beneficial actions. Corollary 2 has important implications for these systems. When a regulation is designed to generate optimal action based on costs and benefits, the best possible response to any future state and its implications should be considered. For example, if some pollution is likely to result from an activity, then regulations should make it relatively easier to pollute those sites that could benefit from destination-driven restoration.

The geographic concentration of waste disposal or of polluting industries in industrial parks is in the spirit of this type of regulatory response. Zoning or other

[13] It may also be desirable under IMC, due to the decreased marginal cost of restoring larger injuries, but the zero marginal cost with ITC makes it more stark.

regulations that concentrate industrial facilities reduce the cost of the externalities created by their pollution. This suggests a violation of the convexity assumptions on the benefit side. But, in addition, it takes advantage of ITC. Instead of having to restore the soil quality so it meets current contamination standards every time it drops below q_e, only to pollute and clean it again, the pollution is concentrated and must be cleaned up only when (and if) it reaches a higher level of contamination. If the site use creates contamination and does not demand immediate cleanup (e.g., dirty industry) it may be efficient to clean it up only when the area is converted to some more sensitive use (e.g., housing) [2].

Liability-Based Harm Prevention

The liability system provides incentive-based environmental regulation, and legally defensible property rights are often the primary source of regulation, especially for major injuries [10]. Such regulation takes the form of lawsuits brought by the holders of regulatory property rights, or their agents, to protect the benefit they receive from the resources. As is generally recognized for tort law in the law and economics literature, optimal incentives for avoiding potentially injurious actions, whether the injury is certain or probabilistic, result from having appropriate liability rules.[14] The responsible party (RP) should be required to pay the net cost it inflicts, that is, the restoration costs plus any residual reduction in the value derived from the injured resource.[15] To get both optimal incentives and optimal remediation, charges to injuring parties and the amount spent on remediation cannot be equated. Typically, optimality is described in terms of the RP paying damage awards that exceed the efficient expenditure on restoration when injury remains that is not socially efficient to restore.

Destination-driven projects, which offer the possibility of a net improvement in quality after injury and restoration, create a new twist. Efficiency requires that the RP pay for the net loss in quality, so accordingly, if quality is increased then the RP should receive a positive credit. That is, the damage award calculation should subtract the value of the net quality from the amount spent on remediation.[16] Such a rule could be justified by arguments about fairness to the RP or to avoid perverse incentives for holders of property rights.[17] The primary concern, however, is that

[14] Optimal incentives are extensively analyzed in the Pigouvian tradition of controlling externalities [1] and its application to tort law in the literature on law and economics [4, Chaps. 8 and 9].

[15] The responsible party, as it is used in the context of Superfund, is often the polluter. However, the RP may not actually take the action that causes pollution, but bear responsibility through ownership or some other relationship.

[16] This principle has implications for calculating prices outside the realm of liability. For example, the city of Boston could respond to a hypothetical rise in sea level as a result of global warming by erecting a dike along its current coastline at some large expense. But in an informal assessment, Dutch engineers suggested that it would not be much more expensive to build the wall across the mouth of Boston Harbor, taking advantage of existing islands, and thereby claim hundreds of square miles of new land from the sea, relieving the congestion of the crowded urban area that cannot expand to the east. (This assessment was recounted by Thomas Schelling, personal communication.) Estimates of the net cost of global warming should add in the cost of the long seawall, but it also should subtract the benefit of the newly reclaimed land.

[17] The cynic will observe that right-holders might have an incentive to promote injuries to resources with an initial $q_0 < q_P$ in order to ‘‘profit’’ from the resulting restoration.

such a credit rule is necessary to create efficient incentives for risking or curbing the risk of environmental injuries.

Without such compensation, several classes of socially suboptimal decisions might result. First, if the RP influences the choice of restoration project, it will clearly opt for a cheaper incremental project (if it is legally sufficient) over a more expensive, socially-efficient destination-driven project. Second, the tradeoff between avoiding the risk of pollution and engaging in productive risky activity will be skewed toward the former, since the cost of pollution internalized by the potential RP will be above the actual social cost. Third, and perhaps most importantly, failure to give credit for improvements will result in the wrong mix of risks. For example, if a RP faces the same destination-driven cost to restore any site it may injure, it will have no incentive to avoid injuring a high quality site that society might be most anxious to protect. Without credit for improving the already diminished quality of lower quality sites, the potential RP will have little incentive to preferentially put such sites, rather than more pristine ones, at risk.

This inefficiency is magnified if a RP might become directly liable for previous injuries that it did not cause (as is the case under Superfund). For example, decommissioned military bases are often converted to uses like shipyards, airports, or heavy industry, that do not require the land to be particularly clean (and will pollute it again). But no firm or local government will want to take over bases, which are frequently very polluted, if the liability they face for any subsequent restoration does not consider the original low quality. A liability rule that makes the new owner responsible for previous injury forces the federal government to clean the sites up to an inefficiently high level before transfer, which results in great expense and delays.

Similarly, without proper credit, potential RPs might inefficiently shift risk or injury to avoid letting a site drop below q_{DI}. If dropping below q_{DI} triggers an expensive destination-driven project, the potential RP may prefer imposing risks on higher-quality sites that do not risk crossing the regulatory threshold. Thus, appropriately reducing liability to reflect net cost can actually increase environmental protection in some cases.

Naturally, the net result of setting liability too high in some cases is a net reduction in environmental injury which many people would applaud. But it would be more efficient to achieve such damage reduction through a more general shift in the tradeoff between environmental protection and industrial activity—a shift that raises the implicit value of all natural resources or of high-quality ones preferentially, rather than just raising it for sites of low quality.

These implications and other contrasts between the standard IMC model and the sometimes more accurate ITC model are summarized in Table I.

V. CONCLUSION

Economic analysis and incentive-based environmental regulations can help to optimally balance the costs of injuries against the costs of avoiding or reversing them. Unfortunately, the traditionally posited shapes of marginal benefit and cost curves bear little resemblance to the realities of restoring or replacing an injured site. Over important ranges of actions, remediation costs exhibit considerable economies of scale. In the extreme case, costs are destination-driven; that is, the cost of remediation depends solely on the terminal quality and not the starting quality.

When destination-driven options for restoration are available, the magnitude of optimal restoration is a ‘‘badly-behaved’’ function of the level of injury and the post-injury quality of the resource. The efficient choice of restoration project may depend discontinuously on the quality of a resource following an injury. Moreover, optimal restoration may result in a net improvement in quality.

Regulatory and liability rules must attend to physical realities. In particular, the economic principle that responsible parties should pay the net cost of the harm they create must be extended so that they receive compensation for over-restoration, just as they would be charged for unrestored injury. Only by understanding the actual shapes of cost functions for injury remediation can we craft rules that will foster efficient levels of environmental risk and polices of remediation.

APPENDIX

Proof of Proposition 1. Clearly, for any given element of **D**, there is a project in **E** that weakly dominates it (since it achieves the same quality at a weakly lower cost), so we need only consider elements of **E**. Define $N(s,t) = V(t) - V(s) - c$, the net benefit from the project in **E** with cost c and final quality t, starting from initial quality $s < t$.

Let $s_a < \min(\mathbf{Q})$. Since **E** is closed with respect to c and t, $V(t)$ is monotonically increasing and defined for all t, and $V(s)$ is constant, the set $\{N(s_a, t)\colon\ t > s_a, t \in \mathbf{Q}\}$ is closed, and thus contains a maximum. Define $P \in \mathbf{D}$ as the project that yields that maximum. (If there are multiple projects tied for the maximum, any one can be chosen.)

Now consider a different initial state $s_b < q_P$. Project P must also maximize $N(s_b, t)$. $N(s_a, t) - N(s_b, t) = V(s_b) - V(s_a)$, so the net benefits differ only by a constant, not changing the optimization decision. It may no longer be optimal to undertake *any* destination-driven project, but the most attractive of those projects will still be the same.

For any initial state $s_b > q_P$, P is no longer an option. But in that case no $R \in \mathbf{D}$ with resulting quality $q_R > q_b$ will be attractive, since if it was not worth the marginal cost, $C_{DM}(q_R) - C_{DM}(q_P)$, to move to q_R from s_a rather than to q_P, then it cannot be worth a higher cost, $C_D(q_R)$, to move there from the higher quality, s_b.

Finally, we observe that P cannot result in quality less than q_e, since it would then offer strictly less net benefit than project P', defined as ‘‘do P and then purchase the incremental restoration from q_P to q_e,’’ the latter step of which produces positive net benefit. Notice that P' fits the definition of a destination-driven project (since we can always pay for the projects even if we are already at a higher quality level where they are useless), so is an element of **D**. ∎

Proof of Corollary 1. This follows immediately from the proposition by noting that if multiple projects P_1 and P_2 produce the same utility then we can choose either one of them to be P. They will produce the same net benefit for any s, and thus will be the optimal project for the same set of s. ∎

Comments on Proposition 2. The proposition follows immediately from the definitions and discussion preceding the proposition. The parenthetical in part (1), which allows for the case where q_{DI} does not exist, reflects the possibility that $c(q)$ will approach zero as quality decreases before a destination-driven project becomes optimal. The practical difference between this and the existence of a destination-driven project with $t = q_e$ is trivial.

Proof of Corollary 2. Part (i) of the corollary follows immediately from part (1) of the proposition. The result in part (ii) applies for any injury that takes a resource from quality q_0 to s where $q_e < q_0 < q_P$ and (Case 1) for $q_{DN} < q_e$, $s < q_{DI}$, or (Case 2) for $q_e < q_{DN}$, $s < q_{DN}$ and $q_0 > q_{DN}$. Part (iii) simply says that once a destination-driven project is optimal, no further cost (ignoring pre-restoration costs) accrues from additional injury. A necessary condition for no destination-driven project ever being optimal is $c(q) \to 0$ as $q \to -\infty$. That is, either a destination-driven project will become optimal as quality drops or the marginal cost of further injury approaches zero; as mentioned above, this is not very different from a practical standpoint. ∎

REFERENCES

1. W. J. Baumol, On taxation and the control of externalities, *Amer. Econom. Rev.* **62**, 307–322 (1972).
2. S. Breyer, "Breaking the Vicious Circle: Toward Effective Risk Regulation," Harvard Univer. Press, Cambridge, MA (1993).
3. T. A. Campbell, "Managing Ecological Liability through Use of Restoration/Acquisition Alternatives," manuscript, National Legal Center for the Public Interest, Washington, DC (1994).
4. R. Cooter and T. Ulen, "Law and Economics," Harper Collins, Glenview, IL (1988).
5. R. W. Hahn and R. N. Stavins, Incentive-based environmental regulation: A new era from an old idea?, *Ecology Law Quart.* **18**, 1–42 (1991).
6. G. E. Helfand and J. Rubin, Spreading versus concentrating damages: Environmental policy in the presence of nonconvexities, *J. Environ. Econom. Management* **27**, 84–91 (1994).
7. T. Holden *et al.*, "How to Select Hazardous Waste Treatment Technologies for Soils and Sludges," Noyes Data Corporation, Park Ridge, NJ (1989).
8. T. Nunno *et al.*, "International Technologies for Hazardous Waste Site Cleanup," Noyes Data Corporation, Park Ridge, NJ (1990).
9. P. J. Parks and R. A. Kramer, A policy simulation of the Wetlands Reserve Program, *J. Environ. Econom. Management* **28**, 223–240 (1995).
10. C. V. Phillips, "Assignment of Property Rights as a Mode of Regulation," Ph.D. dissertation, Harvard University, Cambridge, MA (1995).
11. C. V. Phillips and R. J. Zeckhauser, Confronting natural resource damages: The economist's perspective, *in* "Natural Resource Damages: A Legal, Economic, and Policy Analysis" (R. B. Stewart, Ed.), National Legal Center for the Public Interest, Washington, DC (1995).
12. T. Schelling, "Choice and Consequence," Harvard Univer. Press, Cambridge, MA (1984).
13. J. Soesilo and S. R. Wilson, "Site Remediation; Planning and Management," CRC Press, Boca Raton, FL (1997).
14. D. Starrett and R. J. Zeckhauser, Treating external diseconomies—markets or taxes?, *in* "Statistical and Mathematical Aspects of Pollution Problems" (J. W. Pratt, Ed.), Marcel Dekker, New York (1971).
15. R. G. Stillman, Estimating the remediation of hazardous waste sites, *in* "Hazardous Waste Control" (R. A. Selg, Ed.), Marcel Dekker, New York (1993).
16. O. Tahvonen and S. Salo, Nonconvexities in optimal pollution accumulation, *J. Environ. Econom. Management* **31**, 160–177 (1996).
17. T. Tietenberg, "Environmental and Natural Resource Economics," Scott, Foresman and Company, Glenview, IL (1984).
18. A. Tversky and D. Kahneman, Rational choice and the framing of decisions, *in* "Rational Choice" (R. M. Hogarth and M. W. Reder, Eds.), University of Chicago Press, Chicago, IL (1986).
19. A. Tversky and D. Kahneman, Loss aversion in riskless choice: A reference-dependent model, *Quart. J. Econom.* **106**, 1039–1061 (1991).
20. R. D. Vidic and F. G. Phland, Eds., "Innovative Technologies for Site Remediation," American Society of Civil Engineers, New York (1995).