

MANAGEMENT AND STANDARDS FOR PERTURBED ECOSYSTEMS

Myron B Fiering

C. S. Holling

February 1974

Research Reports are publications reporting on the work of the author. Any views or conclusions are those of the author, and do not necessarily reflect those of IIASA.

Management and Standards for Perturbed Ecosystems

Myron B Fiering*

C. S. Holling**

Introduction

This paper is the first in a proposed series of publications which deals with the application of systems procedures and optimization techniques to a class of problems in ecosystem stability and resilience, and to the consistent evaluation and implementation of environmental standards.

For a long time ecology has been an important factor in setting policies and schedules for management of single species populations--e.g. fish, wildlife, insect pests, and forests. But it has only recently been pushed, or has pushed itself, into a position of importance with respect to many larger issues of policy at all levels of government, thereby taking a dual role. First, its descriptive aspects seek better to understand and tentatively to model features of the enormously rich and variegated dynamic interplay amongst living things; second, its prescriptive aspect seeks rationally to influence policy and to adjudicate some of the inevitable conflicts between economic needs and environmental

* Gordon McKay Professor of Engineering and Applied Mathematics, Harvard University; at present research scholar at the International Institute for Applied Systems Analysis, Laxenburg, Austria.

** Professor, Institute of Resource Ecology, University of British Columbia; at present research scholar at the International Institute for Applied Systems Analysis, Laxenburg, Austria.

(ecological) protection. These conflicts do not imply that polluters are necessarily evil and greedy; we are not so naive as to crusade wholly in the name of generations yet unborn and to march off to battle against politicians and captains of industry. Our analysis is directed at those who would study ecological currents and forces, who would ride them piggy-back rather than subdue them to any end, however noble. It is directed at those who would undertake ecological engineering simply because it may be the most cost-effective engineering for certain laudable commercial enterprises, not merely because it is the gentlest, the most humane, the most natural, or (somehow) the most acceptable.

Our objective is not merely to preserve ecosystems (whatever that means) because it is fashionable to do so, but rather to encourage the systems to work with us. Therefore we are concerned principally with ecosystem policy: its definition(s) and scope, its scientific and socio-economic consequences, its implementation, and its expression in the form of environmental standards. We hold that the implementation of environmental standards and the encouragement of ecosystem resilience are, or can readily be, different manifestations of the same process: to live, more or less harmoniously, with and within the ecosphere. We hold, further, that environmental standards must reflect not only our current perceptions on what is "good" or "bad" about a particular combination of ecosystem state variables, but also how the elements of that array will respond in time

to natural interactions and to man's intervention. These must be made to work in concert.

We first express the dynamics of our ecosystems in models which capture and reproduce what we judge to be the critical biological concepts and rules, fully recognizing the enormity of gaps in our formulations. These models then serve as a laboratory for testing technological, institutional, and policy options. An important part of the art of systems analysis is to constitute as "the system" only that portion of the real world which is relevant for the decisions at hand; we take the view that systems analysis is ultimately a prescriptive discipline, and that it passes through descriptive phases as its models (for decision making) are improved. Consequently, while it is tempting to abdicate responsibility for establishment of environmental standards because of an inadequate scientific basis, the fact is that systems analysis, sympathetically used, should enable us better to use these tentative models and sparse data in making consistent and rational standards. Contrary to the maxim, a little knowledge, used wisely, is often better than none at all; it is always better than a little knowledge used foolishly. Standards are not exogenous to our systems but an integral part of our decision process.

Indicators

The objective of this paper is to introduce some concepts for management of environmental systems. Throughout this work we will use "environmental systems" interchangeably with "ecosystems" because we intend, for purposes of this model, to represent the environment by a small number of surrogate or indicator species or functions, and maintenance of environmental quality to be tantamount to maintenance of these within ranges defined by their stability characteristics. The use of indices to collapse the dimensionality of multi-variate problems in ecology and environmental engineering has a long and spotty history. For example, in ecology there have been large multi-variate studies attempting to identify key species, but except in a few highly specific situations these have not been notably successful or generalizable. Current work suggests that indicators might be defined in terms of roles rather than in terms of species; these roles are defined by size, place, and (perhaps) trophic level. Significant changes in roles occur only when the system is significantly stressed, even while there are, or can be, wide temporal and spatial shifts in the species in the absence of stress.

With respect to environmental engineering, indices for pollution of water and air have been constructed and tested in a variety of applications. These have centered on chemical, physical, and biological characteristics of the media, although a few studies have pointed toward an index property which

reflects the cost of pollution abatement for specific purposes. Additional work has centered on housing characteristics and urban air pollution indices as measures of environmental health; some of the best of these have been used as adjuncts in policy analysis, although the methodologies are not sufficiently general to service a wider need.

Resilience and Standards

Holling [3] summarized a range of experience in classifying equilibria of ecological systems. In particular, he introduced the concept of resilience and contrasted it to stability, noting that a population's resilience, or persistence under stress, is a more important characteristic of its ecological viability than a population count at some moment in time. He described the mechanics of six typical cases of possible ecosystem interaction, using two indicator species to represent the ecological interplay, and plotting all results in the phase plane. The species are identified as X and Y, and the six cases are shown in Figure 1.

It would be useful first to put the basic decision model into its ecological context. Two species (or roles), X and Y, are involved in this highly simplified ecosystem; we assume they live in a lake, which is a relatively simple ecosystem with a well-defined boundary. It is important to the decision maker that these species or roles persist, although it is well known that environmental accidents or pollution

episodes of varying severity can occur, resulting in movements across the phase plane; these movements are typically associated with a kill of one or both species. It is our objective to establish a standard for environmental pollution control such that appropriate capital equipment and operating policies can be employed to maintain X and Y at suitable levels. These levels are not necessarily defined by a minimum alone, but perhaps by a maximum as well. An upper limit might be added because large numbers of one organism can lead to a wild oscillation and ultimately to a new and distasteful equilibrium region. Algal blooms are an example because, once triggered, they can lead to deoxygenation of the water as the abundant plant material decomposes.

It is part of the sanitary engineering tradition to prevent these environmental insults, and great sums are devoted to assure within reasonable limits of probability, that potential waste loads for the system are so treated and managed as to meet performance criteria or standards imposed by society.

Consider the consequences in Figure 1A, Stable Equilibrium, for which the initial condition is at some point (X_0, Y_0) or P_0 on the spiral. Given no intervention by man and no stochastic effects, the tendency of the system is to move inward along the spiral-like trajectory, taking steps of varying size in each successive time interval, and in the limit approaching the equilibrium position (X^*, Y^*) .

Stochastic influences derail the process, the size and direction of the random component being a function of location in the phase plane. The distribution of times of passage from P_0 to P^* , or to some set of points near enough to P^* , might be evaluated analytically if the convergence functions are sufficiently tractable; in most cases Monte Carlo techniques are required. But apart from these details, it is clear that systems characterized by Case A will always migrate toward equilibrium so that while long recovery times might be associated with larger displacements from an interior or quasi-equilibrium position in the phase plane, the fact of recovery itself is a virtual certainty.

Cases B and C demonstrate similar convergence. Case B, Stable Limit Cycle, converges from any point in the plane to a closed loop which shows dynamic rather than static equilibrium. Any perturbation from the stable limit cycle produces ecological pressures which ultimately drive the system back to the cyclic locus. If X and Y alone are plotted against time, the time series would show patterns characteristic of sustained oscillation; this is a typical consequence of simple predator-prey behavior. Case C, Stable Node, shows direct convergence, with no looping, from any point in the X-Y plane to the equilibrium point P^* .

Case D, Neutrally Stable Cycles, has the special property that any displacement results in a new, sustained oscillatory time series. This phenomenon has not yet been

identified in real biological systems; indeed, there may not exist any such systems. But the case is included for completeness.

Case E is of great interest for environmental management; we consider here its special implications with respect to water quality control. We have a closed locus from which inward displacements converge to an equilibrium position (X^*, Y^*) or from which outward displacements diverge to some new domain of stability or to extinction of one (or more) species. Of course, small displacements will not necessarily result in these terminal positions because movements in the phase plane contain random components which might jump the trajectories across the boundary in either direction. It is useful to think of a domain of stability as a mesa surrounded by lowlands. A particle moving on the mesa has a non-zero probability of falling off in one step, and the probability varies according to the location of that particle on the plane and according to the size of the step at any time. Once fallen, unlike a woman, a particle can climb back onto the mesa and re-enter the domain of stability; the likelihood of such a move is smaller than that of falling off.

Another possibility is that the particle, having fallen, wanders around the lowland and ultimately encounters another mesa to which it is propelled. In biological terms, the system flips from one domain of stability to another.

This concept is fundamental to an understanding of biological engineering and the imposition of standards. We

go to great expense to prevent undesirable fluctuations (or what we perceive to be such fluctuations) in the populations X and Y, as if uniformity over time were in and of itself a desirable feature. From Figure 1 we see that dynamic equilibrium is an acceptable ecological state, and that ecostasis may be quite an unnatural scheme to impose on the system. If environmental insult does occur, moving our system far from its equilibrium point or stable orbit, our tradition responds by passing or enforcing laws which prevent additional degradation while allowing time for the abused system to recover. But this response makes no effort to adapt to those ecosystems for which recovery is not defined in terms of specific population levels, and certainly takes no explicit account of the vastly different time scales under which different species come again into harmonious balance or near-equilibrium.

Moreover, if the system flips out of its domain of stability, and if the particle is attracted to a new domain, it might happen that this new orbit is "better" in some measurable way, or at least not worse, than the original footprint or mesa. It might also happen that the optimal strategy in the event of such a flip is the acceleration, by selective seeding or killing of organisms, of convergence to the new domain.

Or, if the new domain is perceived to be worse than the old, it might be reasonable to encourage and accelerate

recovery to the old by selective seeding (or killing), thereby assisting the particle to clamber back onto its original hill. But in re-establishing the old and thereby achieving recovery, it is not necessary to fix the particle at some unique point in the plane but most economically to boost it up over the mesa.

We call this concept Ecological Engineering through the utilization of an Environmental Zoo. It takes specific account of the elasticity of natural systems, encourages them (as we do with structural systems) to roll with the punches, to absorb environmental stress by manifesting strain. Thus if ecological recovery can be made sufficiently rapidly (alone or through seeding), or if alternate domains of ecological stability are suitable, it might be best to reduce capital investment in treatment facilities and to substitute an investment comprised of a combination of costs for less vigorous treatment, maintenance and utilization of the Zoo, and other social and economic costs associated with an increased frequency of environmental insult.

It thus becomes clear that a count of populations X and Y alone is inadequate to define the ecological status of even our simplistic 2-species system. It is not enough to know where on the elliptical footprint we are, how we got there, or even the "elevation" of the footprint relative to the surrounding lowland. What is important is some measure of the "distance" to the precipice; all sorts of geometric

interpretations can be evolved. The thickness of the annulus separating a point from the edge is less significant than the permeability of the boundary; the strength of the guard-rail is more critical than the characteristic of the highway.

A simple measure of the resilience of the system is the propensity of any point on or off the footprint to migrate, within one time interval, to some point on the footprint. If we divide the feasible X,Y space into zones, and associate with each a vector of values which defines the transfer probabilities to all other zones on the plane, we can begin to formalize an analytical framework for studying the behavior over time of a particle or a collection of particles. We might, for example, elaborate the single footprint to a system which encompasses several footprints or domains of stability. Each domain and the interstitial lowlands could be subdivided into a number of distinct regions or states, and to each there could be assigned a transfer probability vector. If the conditions of a Markov Process prevail, or nearly so, these jump probabilities completely specify the process. From them we can deduce the steady-state or residence probabilities for every state, the expected time for transfer between any two states, and other related aspects of system behavior. We return to this formalism in a subsequent section.

Case F, Unstable Equilibrium, is a degenerate case of E for which the footprint is completely collapsed and all points in the X-Y plane lead ultimately to extinction of X, or Y, or

both. This case requires continuous monitoring and management, although it is not obvious when the interventions should be scheduled to obtain optimal control over the ecosystem. For example, if killing is cheaper than seeding, it might be advantageous to wait until the system is in a state northeast of the acceptable region and then reduce X and Y simultaneously, with the next intervention scheduled at a time dependent on the realized rate of movement to the northeast and the relative costs of killing X and Y.

We turn now to a simple model for managing ecosystems which obey the mechanics specified in cases A through F.

A Simple Optimization Scheme for Ecosystem Management

Consider an ecosystem characterized by Case A, Stable Equilibrium. It is evident that A represents a special footprint of Case E in that it occupies the entire X-Y plane. Along each of the infinity of spiral paths we mark steps which are traversed in equal time intervals--say a season or a year or whatever. These are based on deterministic system dynamics, a condition which we will presently relax. In Figure 2 we see several convergent paths, for which the biological evidence ordains that the steps get smaller as the spirals tighten. Particles approach the equilibrium point, reaching it only in the limit, unless a random perturbation jolts a particle directly onto its equilibrium point; such a particle is not forever trapped at equilibrium but will, in

due course, sustain a random shift to some other point in the plane.

Thus it makes little sense to talk of the time required to pass to equilibrium because equilibrium is admissible only as a limiting process. Instead, we draw an envelope around the equilibrium point, and measure time-to-convergence from any point in the plane as the time required to traverse its subtended spiral and intersect the envelope. There is no requirement that the envelope be circular or even symmetric; the principal condition is that it enclose a set of points which can realistically be attained by the system and which delimit "acceptable" combinations of X and Y. In other words, points within the envelope "satisfice."

The spirals and the equal time steps can be mapped into contours of equal time to convergence, where convergence is defined as intersection between a spiral and the locus of admissibility or the target area. The target is not defined by biological or ecological considerations but rather by the boundary of acceptability imposed by the interaction of dynamic social, political, and economic forces. Some targets are large, some are narrow; in a number of cases (e.g., the New Guinea pig culture, Rappoport [4]) it is possible to show how social mechanisms dynamically maintain a wide acceptability domain by forcing a fluctuation of key resources. In any event, convergence is shown in Figure 3 for a set of deterministic spiral trajectories.

Stochastic components can be accommodated by the use of Monte Carlo techniques. Associated with any point in the phase plane is a vector or density whose elements are the probabilities of shifting to the target area within 1,2,3,... time intervals. Thus they are not probabilities of transfer to specific regions of the plane, but rather of transfer onto contours T_1, T_2, T_3, \dots . From these densities, a Monte Carlo analysis can generate contours of equal expected time to quasi-convergence and maps of the probability of rapid (say, 1- or 2- interval) recovery.

Consider a point (x,y) in the X-Y plane shown in Figure 3. This point might be reached after a particular ecological accident--say an oil spill, a thermal discharge, a failure of a municipal sewage treatment plant, a storm surge, etc. The standard for the ecosystem, say a lake, is that management procedures shall be instituted so as to assure that: (i) an accident of this severity occurs no more frequently than one year in $n_{x,y}$, and (ii) when it does occur, the time to quasi-recovery shall not exceed T^* time intervals.

We assume for the moment that no more than one accident can occur in any year. This simplifies the explanations and the example, but in no way invalidates the application of this concept to a more general, realistic case in which several accidents per time period can occur.

The first portion of the standard assures the construction of treatment works or the implementation of process changes at

a scale adequate to reduce to $n_{x,y}^{-1}$ the probability of an accidental shift to (x,y) . The least-cost facility which accomplishes this task has a total cost (capital plus discounted annual OMR) of $K_{x,y}$, whose value depends on the target (a matter of judgment and the political process), the values (x,y) , and the acceptable probability $n_{x,y}^{-1}$. If we examine all minima $K_{x,y}$ on the plane, the facility actually constructed is the most costly or governing one; its cost is

$$C = \max_{x,y} K_{x,y} ,$$

whereupon the probability of an accident at any level (x,y) is changed to $p_{x,y}$, for which $p_{x,y} \leq n_{x,y}^{-1}$. We can put an index on C to identify the policy (or standard) from which it is derived; it is convenient to consider a sequence of policies and to label the corresponding costs C_1, C_2, C_3, \dots .

The second condition of the standard governs recovery time, where we take recovery to encompass all X-Y combinations which lie within the target region (the core or shaded portion of Figure 2). When an accident shifts the system from a satisfactory state to (x,y) outside the locus, the Environmental Zoo is drawn upon to seed the system with x and y organisms (where x and y may be negative) such that the cost of recovery within T^* intervals is minimized.

If costs are linear with seeding levels, if prices of X- and Y-organisms are equal, and if it costs as much to seed as to kill, it is a simple matter to strike a normal to the T^*

contour from the point (x,y) and to read the total seeding cost

$$S(x,y,T^*) = c[|\Delta x| + |\Delta y|] \quad .$$

If unit costs of seeding with X- and Y-organisms are c_x and c_y , respectively, then the normal must be drawn on a plane whose axes are distorted in the ratio $c_x:c_y$. For example, if $c_x = 2c_y$, large changes in X should be discouraged in favor of changes in Y, whereupon the X-axis should be stretched by a factor of 2. This has the effect of changing Y more, and X less, than previously (with $c_x = c_y$).

If unit killing costs are denoted \bar{c} , generally different than seeding costs c , then

$$S(x,y,T^*) = c_x(\Delta x)^+ + c_y(\Delta y)^+ - \bar{c}_x(\Delta x)^- - \bar{c}_y(\Delta y)^- \quad , \quad (1)$$

where terms marked $^+$ are used if seeding is required (Δx or $\Delta y \geq 0$) and those marked $^-$ are used if killing is required (Δx or $\Delta y < 0$) to reach the T^* contour. And finally, if there are economies of scale associated with spreading chemicals or organisms throughout the lake system, the geometric interpretation does not readily apply for reaching the least-cost trajectory. Under these circumstances, a trivial algorithm might have to be incorporated to calculate the least-cost path; in any case, this value is the repair cost associated with meeting the second part of the standard. It is written $R(x,y,T^*)$. The expected annual cost of repair is

$$R = \sum_{\forall x} \sum_{\forall y} P_{x,y} R(x,y,T^*) , \quad (2)$$

and it is reasonable to attach a subscript k to identify the policy under investigation. The total cost of the k^{th} solution is then

$$Q_k = C_k + R_k . \quad (3)$$

Note that changing the standard could involve changing one or more of the following parameters:

- 1) the configuration of the target region, which affects the location of the T^* -contour and hence the evaluation of all recovery costs;
- 2) the value of T^* ;
- 3) the allowable probabilities of deviation from the target domain, which affect the capital costs C_k and the expected repair costs R_k for the k^{th} policy.

The effect of changing these parameters can readily be calculated, and can be presented in several ways. Our example, based on values read from Figure 3, uses the assumptions that killing costs equal seeding costs, or that $c_x = c_y$, so that no distortion of the X-Y coordinate system is necessary. Equilibrium of the ecological system is assumed to obtain at $X^* = 7$, $Y^* = 10$, and the target region is a unit circle centered on this equilibrium, at (7,10). Costs are directly proportional to the sum of movements along the coordinate directions (the Manhattan distance), and the repair costs

$R(x,y,T^*)$ can be plotted as a set of contours, Figure 4. These contours are unique for a given policy which, in this initial example, is simplified by centering the target at the equilibrium point. But, for a more general treatment, it is essential to free the centroid of the acceptability domain or target from its identity with the ecological equilibrium in order to demonstrate the consequences of setting standards independent of the characteristics of the ecological system, and thereby more reflective of socio-economic factors, trade-offs, and aesthetic perceptions.

Suppose, for this more general treatment, we specify that the target area is a circle. In subsequent studies it would be interesting systematically to consider other shapes, configurations, and positions of the target area. The relation between these dimensions and the behavior of the natural ecological system determine the true cost of imposing and maintaining an environmental standard. For example, if we center a circular target area far from the ecological equilibrium, the costs will not only consist of the initial seeding or killing costs required to move the system within T^* time units of the target, but also of the cost of continually maintaining the system, against its inclination, in an inhospitable or abnormal region of the phase plane.

Of course, other policy parameters are subject to tradeoff at the margin. In particular, because the probabilities superimposed on the ecological outcomes which fill the X-Y plane

can be changed by investment, the strategies of prevention and repair can be blended and traded. If we undertake seriously to reconcile different viewpoints, for example the preventers versus the repairers, a panoply of decision making tools is available. These include Paretian Analysis, multi-attribute decision analysis, group decision making, etc. Details of these procedures appear in the standard literature on decision theory, with especially relevant summaries in Dorfman and Dorfman [1] and Dorfman et al. [2]. We will elaborate some of these views below.

Contours could be drawn, and policies compared, for other cases in our set. For Case B, Stable Limit Cycle, the target region could be the same enclosed locus described above or, alternately, an annular ring whose thickness is a reflection of the standard for managing the ecosystem. Such a choice would represent an explicit attempt to work with an intrinsic rhythm inherent in the system, at least to the degree that it appears to be acceptable for social and other reasons. The thickness of the ring, which generally is variable around the locus, is analogous to the size (however measured) of the centrally located target region adopted in Case A. The consequence is that iso-cost contours for recovery under Case B show a curved valley shaped much like the annular target. The implication here is that it would be very expensive to impose an environmental standard which forces the system to the hills locked inside (or surrounded by)

the annulus as compared to relaxing a bit and allowing the system to traverse its natural valley. In other words, the standard should be in harmony with the mechanics of the system and should reflect the economic realities of forcing state variables into uncomfortable combinations which they cannot maintain unless abetted by frequent intervention and management. Similar results can be obtained for the other cases in our basic model system.

Markovian Behavior

The varying step sizes shown in Figure 2 suggest that a given particle spends much more time closer to P^* than at the fringes of the spirals. If we assume that the stochastic process which governs convergence is ergodic, we can also say that at any given moment more particles are located near the center of the galaxy than at the edges. Indeed, the residence probabilities are the relative particle densities at any time. It is reasonable to extend this argument first by dividing the domain of stability into several sub-regions which define the states of a Markov Process. For Case A the domain of stability is the entire first quadrant of the X-Y plane, while for Case E, characterized by a Domain of Attraction, the domain of stability is a sub-space of this quadrant (or a footprint thereon). The probabilities selected in the example which follows are more nearly representative of the Domain of Attraction, Case E, although

it is clear that slightly different entries could be presumed to derive for Stable Equilibrium, Case A. Figure 5 shows four such states in the domain of stability, for which we choose the following sample Markov matrix, which expresses the probability of moving from any state or subregion to any other in a unit time. The selection of four states is a

| | | State, time t + 1 | | | |
|------------------|---|-------------------|-----|-----|----|
| | | 1 | 2 | 3 | 4 |
| State, time t | 1 | .5 | .5 | 0 | 0 |
| | 2 | .5 | .45 | .05 | 0 |
| | 3 | 0 | .3 | .6 | .1 |
| | 4 | 0 | 0 | .7 | .3 |

matter of convenience, but the actual delineation of state boundaries in the phase plane is, ultimately, to be based on careful analysis of the deterministic and random processes which govern particle movement in the domain of stability.

Neither the location of the footprint nor the transfer probabilities in the Markov matrix are based on real data, but are presented only to exemplify the calculation.

In a single term interval a particle cannot migrate over a distance of more than one state in any direction.

States S_1 and S_2 , in the center of the footprint, are occupied much more densely than S_3 and S_4 ; we can visualize an escarpment separating S_2 and S_3 . These states have no specific connection to environmental or pollution standards imposed by the community; they are merely convenient representations of biological evidence, and as such are subject to errors in judgment, differences in opinion, etc. However, it is part of our thesis that the interpretive and prescriptive guidelines to be drawn from models of this sort do not require precise specification of states and of their transfer probabilities in order that prudent policies be recommended.

The steady state or residence probabilities are derived directly from the Markov matrix. Because the Markov Process is defined by a tri-diagonal matrix (a condition which exists when migration is limited to one state in either direction), calculation of the steady-state probabilities is particularly simple and can often be done without recourse to computing machinery. We have $\pi = [0.457, 0.457, 0.076, 0.010]$, suggesting a highly stable regime. Of the four states, S_1 and S_2 are sufficiently close to the target; S_3 and S_4 are not. In other words, we deal with two kinds of "acceptability" in this model. The first, having to do with social and economic issues, delimits the target area, the convergence time, and the frequency or probability of violations of some institutionally imposed standard. This has already been discussed at length. The second kind of acceptability treats the

ecological factors as contained in the phase plane representation, and is independent of human preferences. Thus, of the four states in our Markov model, the natural ecological system will normally lie within S_1 and S_2 , called here the normal subspace, and rarely lie in S_3 and S_4 , called here the abnormal subspace.

Now we examine the impact of a standard. Suppose the community enacts an ordinance, or imposes a standard, which requires treatment of environmental wastes to levels such that the X-Y population counts are constrained to the shaded circular area of Figure 5 most of the time (say at least 90%), with excursions into the attached conical areas not to occur more frequently than (say) 10% of the time. These two regions are called the primary and secondary regions of the standard; we identify the standard by the symbol Z_k , with the index k running across all proposed standards.

This standard is a formal extension of performance criteria imposed on the capital equipment and operating policies of the waste treatment program in the community. It is not necessarily a biological standard, but it must be so expressed as to be mappable into an assignment of allowable probabilities or residence times for each portion of the X-Y plane. While the example in the figure contains two permissible zones, a more general standard can have any number of zones, to each of which there is attached an acceptable residence time.

The location, size, shape, and orientation of the several zones depend not only on the local policy of environmental management but also on the mechanism of the pollution itself. For example, if the source of pollution and its treatment are such that organisms X and Y are jointly encouraged or discouraged, the lower (secondary, tertiary, etc.) zones should be aligned with a positive slope, as in Figure 5. If the organisms respond at the expense of each other, so that they are negatively correlated, the lower zones should be located northwest and southeast of the primary zone.

We note from the figure that 40% of the primary zone is in state S_1 and 60% is in S_2 , while 30% of both secondary zones lies in S_1 and 70% in S_2 . If the standard Z_1 is exactly maintained so that 90% of all outcomes reside in the primary zone and 10% in the secondary zones, the probability that the ecosystem flips from the normal subspace to the abnormal subspace under the influence of policy Z_1 is

$$\begin{aligned}\beta_1 &= 0.9(0.4 \times 0 + 0.6 \times 0.05) + 0.1(0.3 \times 0 + 0.7 \times 0.05) \\ &= 0.0305 \quad .\end{aligned}$$

The calculation is evidently trivial; it represents the cascade of transfer probabilities derived from the Markov matrix and priorities defined by the superimposed environmental standard. It can clearly be generalized to more zones, more states, and more species.

Note that we divide the X-Y space in several distinct ways:

- i) cases the different modes of biological resilience and persistence;
- ii) target the area identified by the community as being socially and politically acceptable;
- iii) subspace normal or abnormal, as defined by the steady-state probabilities deduced from ecological patterns captured by the Markov matrix appropriate to the system;
- iv) states the subregions associated with the system's Markov matrix and residence probabilities, according to which some points on the plane recover rapidly enough to be biologically normal while others do not; and
- v) zones the subregions associated with various treatment options and deduced from the performance characteristics of treatment facilities and policies.

The juxtaposition of all these measures produces a dimensionless scalar, the probability that the system, if all the rules are obeyed, will move within one time period into an

abnormal biological condition. This scalar reflects many characteristics of the earlier measures introduced in our discussion of optimal trajectories for utilization of the Environmental Zoo; these include the size, shape, and orientation of the normal region, as well as the "closeness" of the standard to the edge of the normal region (at which point the system flips into an abnormal or different domain).

However, the probability of movement into an abnormal condition does not, in and of itself, convey enough information because it gives no indication of whether the standard helps the system's natural resilience or works against it. For example, our value of 0.0305 says little in any absolute sense. We cannot deduce if the ecosystem is much better off for the standard. Therefore we introduce some indication of how frequently the system would reside in an abnormal state if no standard were imposed, and from this we gain a measure of increased biological performance attributed to the standard. We define for policy Z_k the parameter

$$\lambda_k = \frac{\beta_k}{\sum_j \pi_j} \quad , \quad (4)$$

where the subscript k represents the (sequential) policy under consideration and the sum in the denominator is taken over all

steady state probabilities which lie in the abnormal portion of the domain. In our example, for which $\pi = [0.457, 0.457, 0.076, 0.010]$, $\sum_j \pi_j = 0.086$, and

$$\lambda_1 = 0.0305/0.086 = 0.355 \quad .$$

A perfect standard always maintains the system in state S_1 so that $\beta_k = \lambda_k = 0$. If the system cannot possibly reach an abnormal position because its biological defense mechanism simply will not allow it, $\sum_j \pi_j = 0$, so that $\lambda_k \rightarrow \infty$ which suggests that the imposition of a standard is foolish (unless there is some other reason for its establishment). Thus the value of λ_k^{-1} measures the extent to which the standard Z_k improves system performance. A value $\lambda_k = 1$ indicates that the standard Z_k neither helps nor despoils the system with respect to maintaining biological acceptability; it is a neutral standard.

Moreover, each standard Z_k is associated with a cost of remedial action should the system, in obeying the standard, actually flip into an abnormal region of the X-Y plane. This is analogous to the repair cost R_k introduced in an earlier section. For Z_1 (Figure 5) the cost R_1 is zero because if the standard is obeyed so that all outcomes reside in the primary and secondary zones, no seeding or killing is required to bring the system to the acceptable level within the requisite time span. However, in Figure 6 we set standards

Z_2 and Z_3 , whose characteristics are significantly different.

Z_2 resides wholly in S_2 so that

$$\lambda_2 = 1.0 \times 0.05 / 0.086 = 0.581$$

and $R_2 = 0$. However, for Z_3 the primary zone subtends a portion of the abnormal region, as does one of the two secondary zones of the standard. The relevant information can be tabulated as follows:

| Zone | Area | Conditional Probability of Being in State | | | |
|--------------|------|---|-----|-----|-----|
| | | 1 | 2 | 3 | 4 |
| Primary | 0.8 | 0.2 | 0.5 | 0.3 | 0 |
| Secondary I | 0.1 | 0.3 | 0.7 | 0 | 0 |
| Secondary II | 0.1 | 0 | 0 | 0.4 | 0.6 |

Thus, of the secondary region marked I, 30% lies in S_1 and 70% in S_2 , etc. The probability of terminating in S_3 or S_4 is the following sum of products:

$$\begin{aligned} \text{start in Primary} & \quad 0.8 [0.2 \times 0 + 0.5 \times 0.05 + 0.3(0.6 + 0.1)] = 0.1880 \\ \text{start in Secondary I} & \quad 0.1 [0.3 \times 0 + 0.7 \times 0.05] = 0.0035 \\ \text{start in Secondary II} & \quad 0.1 [0.4 \times (0.6 + 0.1) + 0.6 \times 1.0] = 0.0880 \end{aligned}$$

from which $\beta_3 = 0.2795$. The sum $(0.6 + 0.1)$ is the probability that a particle in S_3 will remain unsatisfactory by staying in S_3 or moving to S_4 , both of which completely specify the unsatisfactory portion of the plane. We have

$$\lambda_3 = 0.2795/0.086 = 3.25 \quad ,$$

and there is a repair cost $R_3 > 0$ because some allowable (i.e. legal) outcomes are biologically unacceptable!

The cost R_3 is easily calculated. In Figure 6, 30% of the primary zone and all of the secondary (II) zone for Z_3 lie in the abnormal region. If the cost of repair from every abnormal point is calculated as described above, taking account wherever possible of symmetries in killing and seeding costs, etc., and if these costs are multiplied by the probabilities of arriving at each combination, the sum gives the expected repair cost R_k . In the simplest case, for which all three of our cost conditions are met ($c_x = c_y$, killing costs equal seeding costs, and strict proportionality), R_k is proportional to the Manhattan distance between the centroid of each zone in the abnormal region and the point of normalcy. In any case, a suitable numerical approximation to the exact expected cost can readily be constructed.

It is clear that repair costs bear a general relation to the parameter λ_k . Figure 7 shows how λ_k and R_k might be related, with $R_k^*(\lambda_k)$ plotted as a function of λ_k . Many policies

can share the same λ -value, so that each λ is associated with an array of costs. However, it would be futile to adopt any policy which provides a given λ -value at a cost greater than the corresponding $R_k^*(\lambda_k)$, or minimal cost required to meet λ_k . Clearly, small values of λ are associated with zero-cost levels; it is only when a policy lies, in part, within the abnormal region that expected costs are incurred. This is independent of the value of β_k , so that it is not possible to define a general form for the threshold level of λ at which R departs from zero.

Moving Toward a Decision

Figure 7 can be used to initiate decision making; if the decision maker has a clear idea of how he would trade money for environmental (or biological) security, an optimal investment can be directly determined. However, the decision cannot usually be made so cleanly because individuals have ill-defined utility functions and, perhaps more importantly, because environmental issues arouse intense passions which bring vastly different utility and objective functions to the decision making process.

There is a vast literature on this class of problems, and we cannot undertake here to review it. Instead, we refer to a similar model developed by Dorfman and Jacoby (in Dorfman and Dorfman [1] and Dorfman et al. [2]) for dealing with the establishment of a water quality standard

for a stream which flows past an industrial polluter and then through a park area. The conflict is familiar enough: the plant operator wants a relaxed standard, the conservationists and municipal recreation authorities prefer a rigorous one, and the responsible administrator does not know how to resolve the issue. Dorfman and Jacoby recommend a form of Paretian Analysis which allows each participant, or adversary, to evaluate the several options in his own terms, independent of value judgments imposed by other participants or by the central administrative authority. It may happen that certain options are dominated by others; these should be discarded because all parties, by selecting at least one other option or standard Z_k , can be better off or at least equally so.

For those Z_k which are not dominated and which survive other tests of political feasibility, the adversaries are assigned weighting factors, systematically or at random, and the resulting problem degenerates to one of solving for the policy which, under optimal system operation, maximizes the weighted sum of net benefits perceived to accrue to each participant. It can be shown that the solution to this auxiliary problem, maximization of the weighted sum of net benefits, lies on the Paretian Frontier, whereupon it is useful to inquire as to how far the weights must be displaced in order to change certain decisions. This suggests that exact evaluation of the weights might not be required, but

rather than some measure of the solution's sensitivity to selection of the weights might suffice. Other refinements in the technique include the role of uncertainty in selection of an optimal system policy, and current work in this area should be applied to our ecosystem models.

This laissez-faire attitude toward the construction of a net-benefit function for each participant in the decision allows the following classes of costs to be included in the calculation:

- 1) repair costs, to bring the system back to the acceptable region;
- 2) auxiliary source of the resource during the time the system is unacceptable (the cost of doing nothing);
- 3) cost of pollution prevention, or of reducing the probability of an environmental accident;
- 4) aesthetic cost of the pollution episodes;
- 5) economic cost of the imposition of a standard which might drive marginally profitable producers out of production; and
- 6) cost of implementing and monitoring any management program.

In addition, any application of this decision process would require a clear definition of the matrix of probabilities of environmental insult so as to encompass realistic combinations of cascaded events, or of new events occurring before

prior ones are repaired. Indeed, it can be argued that almost every stage of the proposed decision process requires careful examination, tightening, articulation of assumptions, testing, and calibrating; this is the fate of every effort to make prescriptive mathematical models which migrate across the boundaries of disciplines, and is perfectly appropriate. Such detailed examination will be the subject of additional documents in this series.

Other Cases

Much of the preceding discussion is based on Case A, Stable Equilibrium, or Case E, Domain of Attraction; here we consider the dynamics of some of the other cases.

A 5 x 5 Markov matrix is chosen for Case B, Stable Limit Cycle. State S_3 , an annulus around the limit cycle, represents the normal region, with S_1 , S_2 , S_4 , and S_5 circumscribed as shown in Figure 8. A sample matrix is

| | | State | | | | |
|-------|---|-------|-----|-----|-----|-----|
| | | 1 | 2 | 3 | 4 | 5 |
| State | 1 | 0.4 | 0.6 | 0 | 0 | 0 |
| | 2 | 0.1 | 0.3 | 0.5 | 0.1 | 0 |
| | 3 | 0 | 0.1 | 0.8 | 0.1 | 0 |
| | 4 | 0 | 0.1 | 0.6 | 0.2 | 0.1 |
| | 5 | 0 | 0 | 0 | 0.7 | 0.3 |

Note that particles can jump through more than one state; this does not reflect an important methodological change but rather a simple numerical addendum. The steady state probabilities are $\pi = (0.023, 0.137, 0.703, 0.120, 0.017)$ and

$$\sum_j \pi_j = 0.297.$$

^j Figure 8 also shows a standard Z_4 , for which

$$\begin{aligned} \beta_4 &= 0.1(0.5) + 0.1[0.6(0.5) + 0.4(1)] + 0.8[0.7(0.5) + 0.3(1)] \\ &= 0.64 \end{aligned}$$

This calculation is clear enough; it should be noted only that 60% of the northeasterly secondary zone lies in S_2 (with a probability of 0.5 that the next move will take it out of S_3) and 40% lies in S_1 (which cannot jump to S_3), and similarly for the primary zone. We have $\lambda_4 = 0.64/0.297 = 2.15$, and $R_4 > 0$ because some of the standard overlaps the unacceptable zone. In fact, for the case of a system characterized by a Stable Limit Cycle, a standard shaped like Z_4 and oriented in the center of the cycle is not useful because it concentrates outcomes in the central location, thereby failing to account for the centripetal tug toward a dynamic equilibrium represented by the orbital path. To be sure, the system can be restrained, or kept on a tight leash, but only at the expense of continual adjustment and treatment. This is not necessarily bad because there may be perfectly good economic and social reasons for doing so. But in the absence of such justification,

it would appear prudent to allow the system to meander through its cycle and to direct elsewhere those economic resources which might have been required for ecosystem management.

Cases C and D, Stable Node and Neutrally Stable Cycles, are simple extensions of Cases A and B, and merit no further comment here.

Case E, Domain of Attraction, with which we have already dealt, encompasses the notion of expulsion from one domain and attraction to another. There is a small mathematical complication with Markov models of ecosystems which can migrate to extinction, and care must be taken in defining the several states. If we assume that state S_0 is extinction of one or both species, and that there is no way out of S_0 , then the transfer probabilities $p_{00} = 1$ and $p_{01} = p_{02} = \dots = 0$. This is an absorbing process because if a particle reaches S_0 it is trapped; S_0 is called a trapping state, and so long as each particle can somehow reach S_0 , sooner or later all particles will get there and be trapped forever. This cannot be allowed in our model simply because it does not reflect biological reality, whereupon the states must be defined so that extinction is a state embedded within another domain from which it is possible, sooner or later, to migrate to yet another domain or system or to flip back into the original configuration. Our model must encompass the reality which recognizes that biological systems do not shift from one domain

to another and then simply stop to go extinct. They grow, evolve, mutate, adapt. To accommodate these concepts we must generally build larger, more tedious models which require more data and more manipulation; from a conceptual point of view, the extension is trivial. Fortunately, some ecologists feel that there is enough data available to investigate the implications of these complications.

And finally, Case F, Unstable Equilibrium, is a straightforward development based on Case E.

Summary

This paper introduces some new constructs for measuring the effectiveness of environmental standards and for formalizing the economic and social costs of meeting them. It emphasizes the element of recovery time (and of its manipulation) in ecosystem management, and thereby lays the mathematical foundation for further study of an Environmental Zoo and for evaluation of long-term exposure to low-level pollutants. The biological models presented here are taken directly from Holling, whose work distinguishes importantly between the properties of stability and resilience, and who is continuing even now to apply these concepts to management of the spruce budworm in the province of New Brunswick. In place of conclusions, this paper identifies a few potentially fruitful programs for further work.





FIGURE 1. ECOSYSTEM STABILITY (after Holling [3]).



**FIGURE 2. DETERMINISTIC SPIRALS
WITH EQUAL TIME STEPS.**

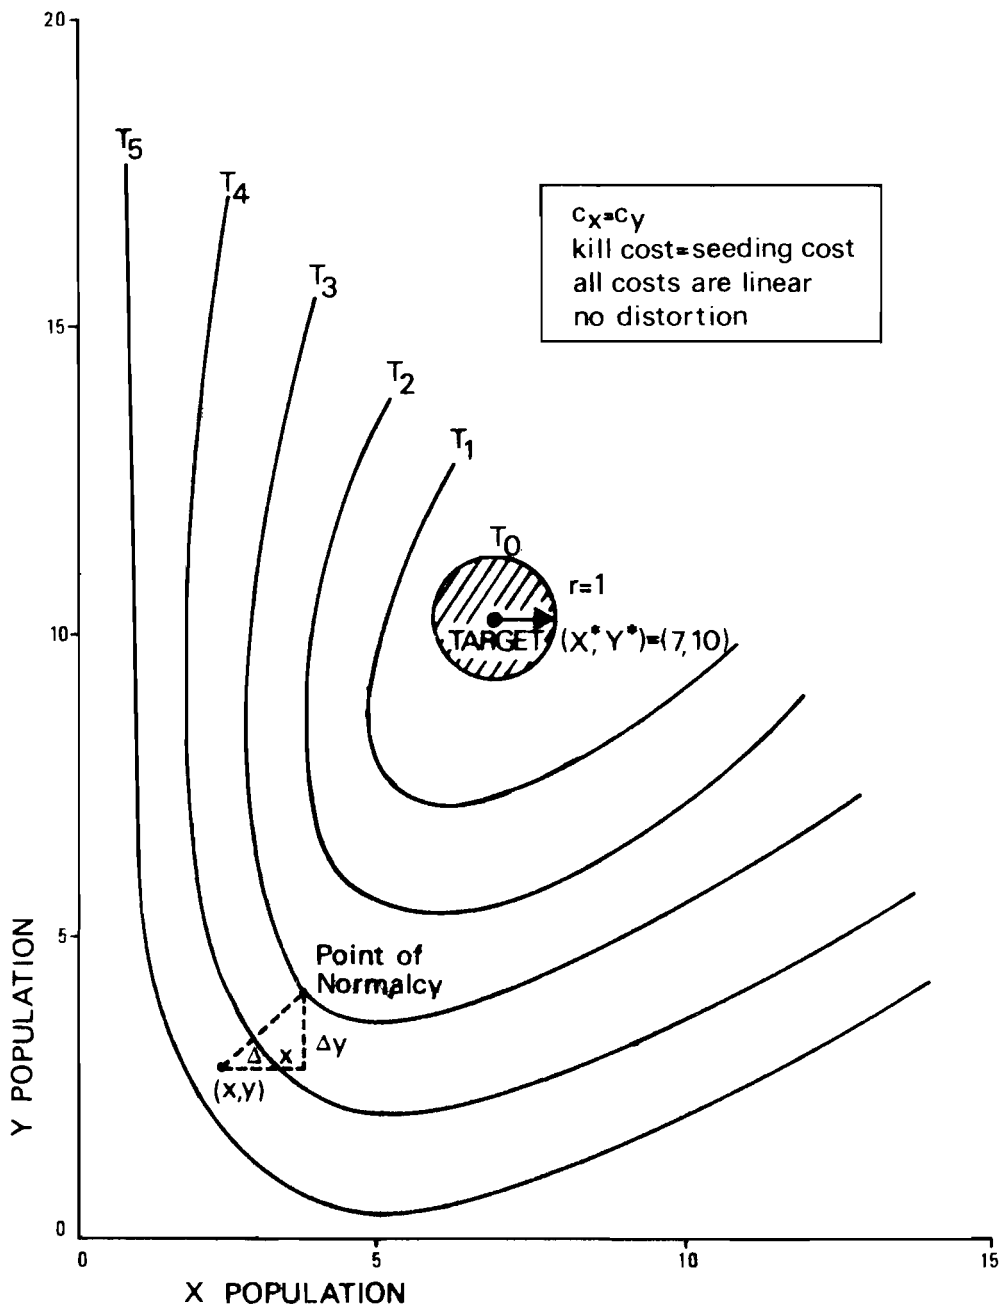


FIGURE 3. CONTOURS OF EQUAL TIME TO RECOVERY.




FIGURE 4. CONTOURS OF EQUAL COST OF REPAIR TO T_3 .




FIGURE 5. EVALUATION OF POLICY Z₁, CASE A OR E.




FIGURE 6. EVALUATION OF POLICIES Z₂ AND Z₃, CASE A OR E.

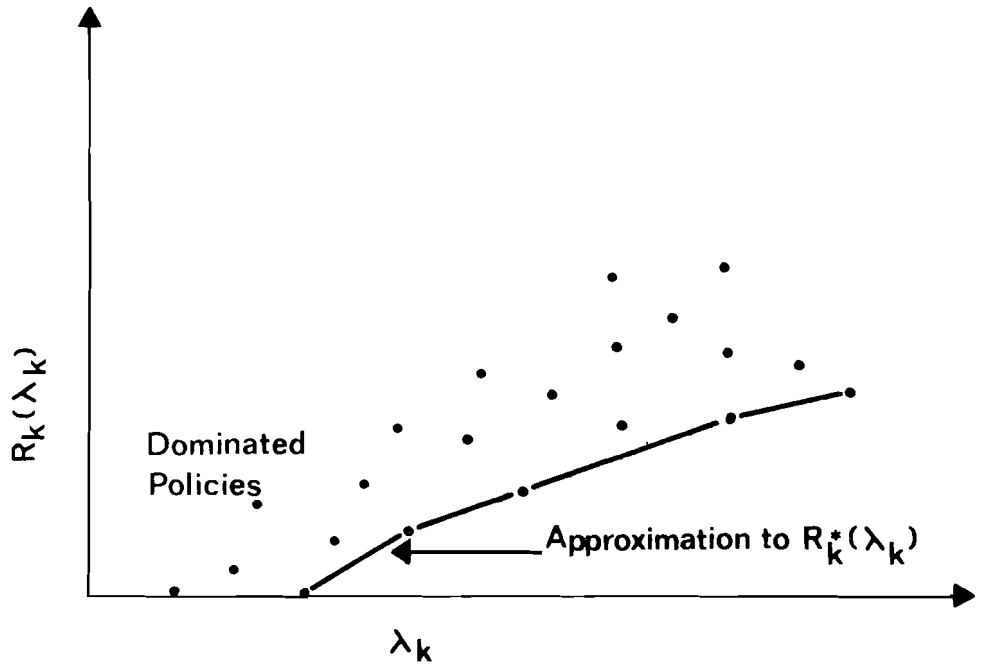


FIGURE 7. MINIMAL COST FOR MAINTAINING λ_k .




FIGURE 8. EVALUATION OF POLICY Z₄, CASE B.

References

- [1] Dorfman, Robert and Dorfman, Nancy S. (ed.). Economics of the Environment. New York, Norton, 1972.
- [2] Dorfman, Robert, Jacoby, Henry D., and Thomas, Harold A., Jr. Models for Managing Regional Water Quality. Cambridge, Mass., Harvard University Press, 1972.
- [3] Holling, C.S. "Resilience and Stability of Ecological Systems." Annual Review of Ecology and Systematics, 4 (1973), 1-23.
- [4] Rappoport, Roy A. "Ritual Regulation of Environmental Relations among a New Guinea People." Ethnology, 6 (1967), 17-30.