Resilience and sustainability: the economic analysis of non-linear dynamic systems

W.A. Brock
K.-G. Mäler
C.A. Perrings

1. Introduction

The discussion of resilience in this book relates mainly to the properties of dynamical natural systems-especially to hydrological and ecological systems. In this chapter we consider the implications of the resilience of such systems for the economics of natural resources. To do this we revisit some of the fundamental assumptions of economic theory. We argue that where it is relevant to discuss the resilience of the system-where it can flip from one state to another if sufficiently perturbed-we need to relax certain of these assumptions. In particular, we need to relax the assumption of the convexity of production sets. This turns out to have far-reaching consequences for the way we think about the economics and management of natural resources, and about the sustainability of resource-based development.

At one level the insights are not new. Economists have been aware of the general problems posed by the non-convexity of preference and production sets for a number of years.1 However, they have paid little attention to the implications of non-convexities for the dynamics of economy-environment systems. Nor have they explored their implications for management and policy. In what follows we discuss both.

The chapter is in six parts. We first identify the way in which basic resource allocation may be affected by non-convexities, and consider the implications of this for the efficiency of the economic process. Following this we explore the role of non-convexities in the dynamics of economy-environment systems and their effect on equity, sustainability and resilience. We also relate this to the variability of natural resources, and to the related problems of uncertainty and irreversibility. We then illustrate our comments through a specific example-shallow lakes subject to nutrient loading, and discuss the implications for . We conclude by discussing the implications of non-convexities for resource management and policy, and for the economics of the environment in general.

2. Convexity and resource allocation mechanisms

The objective of most societies is to maximise the well-being of their citizens. For that we need to be able to represent social well-being in a practical way and to be able to compare the well-being of different individuals. We will start with the impact on human well-being of changes in the quantity or quality of environmental resources. We will therefore focus on individual well-

[1] Since the articles surveyed in Brown (1991).

being.

A utility function is a practical device by which we can represent the preferences of an individual. The utility function is an assignment of real numbers to each bundle of commodities and services (including natural resources or ecological services) available to the individual. Let x and y denote two such bundles. Then if the individual prefers x to y, the utility assignment gives a higher utility for x than for y or

$$u(x) > u(y).$$

Note the order of causality here. Utility is higher from x than from y because the individual would choose x instead of y. In this sense, the utility function is a representation of the preferences of the individual. It is <u>not</u> true that the individual prefers x to y because utility is higher with x than from y. The causality is just the opposite.

A very natural ethical hypothesis is the principle of consumer sovereignty: It holds that individuals are the best judge of their own welfare, given that they have no less information than anyone else about the technical and scientific characteristics of the bundles they are considering. On the basis of this principle it is clear that the utility function can be interpreted as an index of the individual's well-being.2 We do so in what follows.

The two fundamental theorems of welfare economics

By the *Pareto criterion* is meant that one allocation A of resources is socially preferred to another allocation B if at least one individual is better off with A and no one is worse off with A. Note that this is a *partial ordering* as not all allocations can be compared in this way. Two allocations such that some people are better off in the first and others are better off in the the second cannot be compared by applying the Pareto criterion. Let the utility of individual h in allocation A be u^{hA} and the utility in B u^{hB}. Then A is better than B according to the Pareto criterion if $u^{hA} \geq u^{hB}$ with strict inequality for at least one individual.

By *Pareto optimality* is meant an allocation such that it is not possible to improve the

---

2 Let x and y be two bundles of commodities. If an individual prefers x to y we write
$y \underline{\phi} x$.
If he strictly prefers x to y we write
$y \phi x$
and if he is indifferent between the two bundles we write
$y \approx x$.
We want (for practical reasons) a simpler representation of the preferences.
It can be shown (under some reasonable assumptions) that there exist a function u mapping commodity bundles to real numbers such that
$y \phi x$? $u(x) > u(y)$ and
$y \approx x$? $u(x) = u(y)$.
The function u is the utility function of the individual..

If $u(x) > u(y)$, we will also interpret that as meaning that the individual is better off with x than with y. Otherwise they would not prefer x to y! This reflects the assumption of consumer sovereignty. Usually we impose on the utility function various structures, and for doing environmental valuation, special structures are needed.
Note that if we use a utility function to represent the preferences of an individual, the utility function is not unique. If $x \underline{\phi} y$ implies that u(y) ? u(x),then any monotonic increasing function of the utility function is also a utility function. For example, if $u = xy$, then $\ln u = \ln x + \ln y$ or $u^2 = x^2y^2$ are also utility functions representing the same preferences. They are equally valid representations.

situation of one or more individuals without harming other individuals. A Pareto optimal allocation is thus a "maximal" allocation with regard to the Pareto criterion. Note that a given Pareto optimal allocation may not be socially desirable at all because the distribution of well-being, that is the utilities, may be very uneven.

We can now formulate the two basic theorems in welfare economics[3]:

The *first theorem of welfare economics* says:

> If there is a competitive equilibrium such that: all goods can be assigned property rights (and therefore be traded), no individual and no firm can affect the prices, producers are maximising their profits, consumers are maximising their utilities and all markets clear, then the resulting allocation of resources is Pareto optimal.

Note the assumption about property rights. We know that this assumption does not correspond to the real world. It has been studied in depth in many contributions to economic theory, and the importance of the lack of defined property rights is well-understood.[4] The assumption of perfectly competitive markets is also violated in the real world but is probably of less importance than the first assumption.

There is, however, one problem that this theorem does not address, namely whether the Pareto-optimal resource allocation generated by the markets is socially desirable or not. The outcome of the market process may be a distribution of well-being among citizens which is not regarded as good. In fact, the resulting allocation may be highly inequitable. Society may accordingly prefer an allocation that is not Pareto optimal but is more equitable. The second theorem addresses this problem.

The *second theorem of welfare economics* says:

> If production sets are convex and closed, if preferences are convex and continuous and if all goods and services can be assigned property rights, then to each desired Pareto optimal allocation one can find an initial distribution of wealth such that the resulting competitive equilibrium is the desired allocation.

This means that if the assumptions hold, we can always achieve any desired allocation by using the market mechanism and by redistributing the initial wealth among the individuals. Thus we can let the markets do almost everything if we first redistribute the initial wealth in a suitable way. We will in later discuss briefly how we should think of these redistribution and how one should make the trade offs between equity and efficiency.

Note the three conditions:

i) convexity
ii) continuity

[3]Welfare economics is that part of economics that is normative. It is used to make recommendations.
[4]See for example Ostrom (1990), and Ostrom, Gardner, and Walker (1994).

iii) private property rights.

Convexity

Convexity in production means loosely that we have diminishing returns. More strictly we can represent convexity as follows: A set is convex if given two arbitrary points of the set, the line connecting these points belongs to the set. Production is convex if the production set, that is the set of all technically feasible bundles, is convex (see Figures 1 and 2). There are good theoretical and empirical reasons to assume that this is the case in most instances, but as we will see it may not be the case in every instance.

Convex preferences are defined in a slightly different way. Indifference curves are curves connecting bundles that provide the individual with the same level of utility, that is the individual is indifferent between the bundles. Since convexity of preferences is of little concern in this chapter, we will not discuss it any further.

We can now get an intuitive understanding of the economic importance of convexity and the two theorems of welfare economics. Consider an individual who is both producer and consumer. Assume they can produce two goods, x and y. As producer, they have to choose a bundle of the two goods that is feasible in the sense that it belongs to the production set defined in Figure 3 by the transformation curve T-T and the area below it. As consumer they have to choose a bundle that yields the highest attainable level of utility. It is apparent that it is best to produce and consume at the point where the indifference curve is tangent to the transformation curve. At this point, the two curves have a common tangent, P-P. Convexity of the production and preference sets is what guarantees that the indifference and transformation curves can be separated by a straight line. Without convexity, this cannot be guaranteed.

More importantly, the convexity of preference and production sets underpins the efficiency of prices in the allocation of resources. The slope of the tangent may be interpreted as the relative price of y in terms of x. The importance of this is as follows. If the individual as producer maximises profit, they will choose the production bundle that corresponds to the point of tangency between the price line P-P and the transformation curve T-T. If the individual as consumer maximises utility, they will choose the consumption bundle that corresponds to the tangency of the price line P-P and the indifference curve I-I. These points are the same. Therefore, we can decentralise production and consumption decisions whilst still achieving the social optimum. The producer can choose the production bundle knowing only the relative prices of x and y. The consumer can choose the consumption bundle knowing only the same relative prices. This is the intuitive background for the second welfare economics theorem.

Now consider what happens if the production set is not convex. In Figure 4 the production set defined by the curve T-T is not convex. The optimum point for the individual as both producer and consumer is given by point A, where the indifference curve I-I and the transformation curve T-T are tangent to each other. However, at this point, the production and the consumption decisions cannot be ‘separated’ by prices, and the second theorem of welfare economics does not hold. The outcome is not a social optimum. Given the prices defined by the slope of the tangent at A, the producer could make higher profits at some other output combination. At point B, for example, they would make a profit corresponding to the dotted line. Thus, with this non-convexity it is no longer possible to decentralise decisions between consumers and producers and still achieve a social optimum!

Continuity

Continuity is a mathematical device that is usually (but not always) without significant economic meaning, and we will not dwell anymore on this concept. We later consider models of ecological systems that superficially seem to generate discontinuities, but as a matter of fact, they don't. In particular, we will look at a lake that flips between an oligotrophic state and and a eutrophic state and a rangeland that flips between a wooded and a grassy state. It seems that such flips imply a discontinuity but as the flip takes time, it is not mathematically a discontinuity.

Property rights

Without well-defined property rights, markets will not be established for all goods and services and the result is that incentives will be distorted. The reason is that well-defined property rights provide the owner of the rights with an incentive to manage the resource in a socially efficient way. A person who owns an asset, such a piece of land, has an incentive to manage that asset efficiently because they will themselves bear the cost of mismanagement. When property rights are not well-defined, someone else will bear this cost. Given that I do not own the atmosphere. I will not bear the full cost of the environmental damage when I pollute the air. The lack of property rights in such cases leads to externalities-effects that are external to the transactions between individuals. The result will be too much pollution.

Take the case of grazing land. If there were no well defined property rights to grazing lands, for example, there would typically be overgrazing. This is because no single herder would have to consider the full social cost of adding more cattle. If the grazing land were a common property with unlimited access to the members of a particular community, there would still be overgrazing because each member of that community would have an incentive to add cattle beyond the socially optimal level. However, if the common land were managed by some social device that limited the access enjoyed by right-holders, then the grazing land could be managed efficiently. Similarly, if the land were divided into privately owned pieces, each land owner would have an incentive to use their land now in such a way as to be able to use it again in the future. Once again, the grazing land could be managed efficiently. In fact, many environmental problems are due to property rights failures.

There are two basic reasons for property rights failures-policy decisions and costs of establishing property rights. In many countries, environmental resources are regarded as either publicly owned or open to all. For example, land that has not been claimed by anyone is often regarded as open to all. Thus there are incentives to overuse this land. One reason for this may be that the costs of introducing property rights in environmental resources are very high due to the nature of the resource. For example, it is difficult to think of how property rights might be assigned to the global climate. The reason in this case is the public good nature of the resource. A change in climate that affects me will also affect you. Of course, we could think of individual rights to emit green house gases but the climate itself will continue to be a public good.

In a rightly celebrated paper Coase(1960) presented what later became known as the Coase theorem. For our purposes this states that if there are well-defined property rights in some environmental effect, if there are no transaction costs and if wealth effects are negligible, then the outcome of bargaining between the interested parties will be Pareto optimal regardless of the distribution of those rights. Furthermore, lack of agreement about such effects is an indication that the transaction costs are too high to motivate the establishment of a market in the effect. The absence of an agreement is therefore socially optimal. Applied to a simple pollution problem, the Coase theorem says that if those who are damaged by pollution are not able to bribe the polluters to reduce pollution, then it is socially desirable to have that pollution. The gains from pollution control will not outweigh the cost of abatement and the transaction costs.

The Coase theorem has been criticised on many grounds. First, if there are more than two parties involved (and generally there are) the incentives may go completely astray, and the Coase theorem is no longer valid.[5] Second, if there is asymmetric information between the affected parties-some parties know more than others-the outcome of negotiations will not generally be Pareto optimal. However, in this connection the important limitation of the Coase theorem is that the transaction costs, that is the cost of assigning rights and carrying out negotiations between right-holders, will increase substantially when the feasibility sets for the parties are non-convex. The reason is that if these sets are convex, the parties can arrive at the optimal solution by following a step-wise process known as the gradient process. If for example the parties make incremental bids in such a way that the bidder is better off if the bid is accepted, then the process will converge at the efficient outcome. However, without convexity, such a process will generally not converge to the optimum. We consider an example of this later.

The main point of this brief discussion of basic resource allocation mechanisms is that they all require the convexity of production sets in order to operate smoothly. These convexity assumptions are often (but far from always) valid when we limit our analysis to conventional production functions. As we shall see later, however, extending production functions to include non-market environmental inputs, frequently compromises the convexity of the production set.

3 Equity, sustainability and resilience

If we are to judge whether one allocation is better than another we have already made the point that we need to compare outcomes with different implications for the distribution of wealth and income. The only way to judge whether some change increases welfare or not is by explicitly introducing ethical values on distribution of welfare. This is typically done

[5] See Dasgupta and Mäler (1995) for a discussion of the garbage game.

through a social welfare function. Let there be H individuals that are affected by a particular change or policy and let each individual have utility u_h. Let there be a function $W(u_1, ..., u_H)$ which we will suppose is increasing in all its arguments. We will interpret this function as a social welfare function if it represents the preferences of a social "decision maker". We know from Arrow's impossibility theorem that it is impossible to derive the social welfare function from individual preferences in a way that is consistent with some reasonable conditions. However there often does exist information that makes aggregation of individual preferences possible. One can therefore interpret the welfare function either as an aggregate of individual preferences (given certain assumptions on information) or as representing the preferences of a social decision-maker.

Consider a change in society from A to B. This change implies that individual utility levels change from u_h^A to u_h^B, and that social welfare changes accordingly

$$\Delta W = W(u_1^B, ..., u_H^B) - W(u_1^A, ..., u_H^A)$$

If $\Delta W > 0$, then social welfare increases according to the welfare function. This provides one very useful way of thinking about the sustainability of economic and environmental change. There is no generally accepted definition of sustainable development (and it is doubtful whether the concept has scientific validity). The definition offered by the Brundtland Commission-development that meets the needs of present generations without compromising the ability of future generations to meet their own needs (WCED, 1987)-implies something about the change in welfare over time.

Pezzey (1997) identifies three possibilities: that welfare is non-declining over time ($dW/dt > 0$ always); that welfare does not exceed the maximum constant level of welfare ($dW/dt \leq W_{tmax}$ always); or that welfare does not fall below some prescribed minimum ($dW/dt \geq W_s$ always). He defines these as, respectively, sustained, sustainable and survivable development. But for whom is W a social welfare function in these definitions-the present generation or the present generation and generations yet to come? In the discussion on the meaning of sustainable development, many have argued that W(t) should represent the welfare of the generation living at time t. However, this would severely restrict policies. For example, poor countries may need to reduce their current consumption in order to increase the capital stock for future generations. Taking such criteria as non-declining or constant welfare seriously would therefore imply that such countries could never increase the savings rate. A more flexible interpretation is that W(t) represents the welfare of the current and future generations together.

Dasgupta and Mäler (2000) have shown that given some “technical” assumptions, there always exist shadow prices p_i on all reources (including ecological resources) X_i such that an operational criterion for sustainable development is that

$$\Sigma p_i \, dX_i /dt \geq o$$

in the present and in all future time periods. Thus, economies are on a sustainable development path if and only if the sum of the values of changes in capital stocks does not decrease over time. This corresponds to what has been called genuine savings in reports from the World Bank.[6] However, amongst these "technical" conditions guaranteeing the existence of the appropriate shadow prices is the differentiability of the so called value function with respect to the initial endowments of assets, that is of the social welfare function for the present and the future generations. It can be shown that for non-convex feasibility sets, these shadow prices may not exist for all possible initial stocks. Prices may not exist to inform us whether we are on a sustainable path or not.

Resilience and sustainability

An alternative approach is to consider the sustainability of an economy and its supporting environment in terms of its capacity to absorb stress and shock without fundamental change. For any economy there are many possible states, each delivering different levels of welfare to society. In this approach the sustainability of any particular state depends on the properties of the stability domain corresponding to that state. In the ecological studies reported in this volume this is typically analysed in terms of the resilience of the system in each state.

One measure of resilience is the magnitude of disturbance that can be absorbed before a system flips from one state to another (Holling, 1973).[7] Holling (1986) describes ecosystem behaviour in terms of the sequential interaction between four system functions; exploitation or colonization of disturbed ecosystems; conservation as biomass accumulates; creative destruction where an abrupt change caused by external disturbance releases energy and material that have accumulated during the conservation phase; and reorganization where released materials are mobilized to become available for the next exploitive phase. Resilience is measured by the effectiveness of the last two system functions. It is crucial to the ability of the system to satisfy 'predatory' demands for ecological services over time, and to cope with both sustained stress and shock. This measure is the one that is reflected in the papers in this volume.

It has been argued that this measure, and the concept behind it, offers a useful way to address the sustainability not just of ecological systems, but of jointly-determined ecological-economic systems (Common and Perrings, 1992; Levin et al, 1998). Indeed, the approach has implications for the way we think about the dynamics of any stochastic, evolutionary system. Levin et al (1998) argue that sustainability as a concept is more pertinent in stochastic systems away from equilibrium than in deterministic systems at equilibrium.

The link between the resilience of systems and the probability of their collapse or change of state is reflected in the literature on the analysis and management of environmental risk. Although deterministic bioeconomic models for the optimal utilisation of natural resources generate sustainable (steady state) solutions, they necessarily ignore inherent or environmental stochasticity in the modelled relationships, and in trophic or competitive interactions. Randomness has been incorporated into such models via assumptions of

---

[6] The concept was first used by Pearce. and Atkinson (1993).

[7] A second definition of resilience refers to the properties of the system near some stable equilibrium (ie. in the neigbourhood of a stable focus or node). This definition, due to Pimm (1984), takes the resilience of a system to be a measure of the speed of its return to equilibrium following perturbation. The two measures are related

stochasticity in model parameters, random catastrophe or density-dependent risk of collapse (Reed, 1988; Tsur and Zemel, 1994). Density dependent risk of collapse includes both the existence of a density-dependent hazard function (Reed, 1979; 1988) or thresholds which, if reached, trigger the immediate collapse of the stock (Tsur and Zemmel, 1994; 1997). A density-dependent threshold implies that increasing stress on the system raises the probability that it will flip from one state to another, and so corresponds well with a measure of resilience in the sense of Holling. The implications of this for the management of susceptible systems has yet to be considered.

The link between resilience, diversity and risk has been studied in ecology against the backdrop of a long-standing dispute about the relation between the complexity of ecological systems, their diversity and their stability (May 1973; Elton, 1975), and an alternative proposition that diversity supports not stability, but resilience (Holling, 1973; 1986) and ecosystem functioning (Vitousek and Hooper, 1993; McNaughton, 1993). Experimental research of grasslands has now shown that ecosystem productivity increases significantly with plant biodiversity (Tilman et al, 1996). This is because the main limiting nutrient, oil mineral nitrogen, is utililized more effectively the greater the diversity of species. These results have led to the proposition that the sustainability of soil nutrient cycles and so of soil fertility increases with biodiversity.

More generally, the resilience of any ecosystem with respect to variation in environmental conditions depends upon the existence of species capable of supporting the key ecological functions as conditions vary (Perrings, 1995). Deletion of a species important under some conditions will have little effect on ecosystem functioning if there are other species capable of stepping in as substitutes. If there are no substitutes, however, the deletion of some species can trigger a fundamental change from one ecosystem type to another—from forest to grassland, or grassland to a shrubby semi desert, for example (Westoby et al 1989). The importance of the mix or diversity of species for the resilience of ecosystems lies in the fact that species which are 'redundant' in one set of environmental conditions may be critically important in other environmental conditions. Resilience has been shown to depend on the functional diversity of species supporting critical structuring processes (Holling et al, 1995).

The significance of this for resource-based development is that agroecosystems—ecological systems whose species mix is transformed for the purpose of agriculture—may be especially sensitive to species deletion precisely because they are already simplified by the exclusion of competitor or predator species (Conway, 1993). The specialisation gains from simplification of agroecosystems typically involve a reduction in the resilience of the system. The costs of a reduction in resilience include, for example, the cost of the herbicides, pesticides, fertilisers, irrigation and other inputs needed to maintain output in the simplified system. They include the cost of relief where output fails, relocation where soils or water resources have been irreversibly damaged, rehabilitation where damage is reversible and insurance against crop damage by pest or disease. If the system loses resilience and flips from one state to another, they include forgone output under the new state.

There have been few attempts to estimate the impact of changes in relative prices on nonmarketed biological resources, and fewer still that relate these changes to the resilience of agroecosystems. Loss of resilience implies both an increase in the time taken to return to equilibrium following some shock, and a narrowing of the range of environmental conditions over which the system can maintain the flow of ecosystem services. It is economically interesting if: (a) it alters the risks associated with a given set of environmental conditions, and (b) the value or potential productivity of the new and old states is different. If changes of state are either irreversible or only slowly reversible, a change in the potential productivity of the system imposes costs or confers benefits on both present and future users of that system. That is, it has implications for our measure of welfare.8

Resilience, institutions and the evolution of economy-environment systems

Consider the following simplified way of thinking about the evolution of an economy-environment system. Suppose that there are finite number of resources denoted $X_t = (x_{1t}, ..., x_{nt})$. Decision-makers are assumed to allocate resources through actions, a. These actions describe the consumption and production activities of economic agents. They affect the probability that the system in one state will converge on any other state. That is P(a), defines the transition probability between states as a function of the consumption and production activities of decision-makers. The activities of economic agents are, in turn, determined by a set of behavioural rules that depend on the institutional and cultural conditions in society. A familiar example of such a behavioural rule is profit maximisation, which is associated with the decentralised decisions corresponding to the institutional conditions in competitive market economies. As we have already remarked, given relative prices, application of this rule determines the optimal combination of inputs in production and the optimal combination of outputs.

We can call such behavioural rules 'policies', and denote them by u_t. Hence the set of resources evolves according to

$$X_t = F(X_{t-1}, u_t)$$

The policies that guide people's activities are to a large extent determined by the institutional conditions—the 'rules of the game'—within which they are made. That is, institutional conditions determine the logic of optimising behaviour in a way that makes the decision-maker's behaviour fully predictable once the institutional conditions are given. The logic of open access, as we have seen, ensures that decision-makers will choose to use resources up to the point where total revenue and costs are equal, and so on. The implication of this is that for given institutions and we can identify the long-term

8 If the loss of resilience of a managed or impacted ecological system is associated with a change in its long-run productive potential it is in principle observable through its effects on the value of economic output. Perrings and Stern (1999) estimate change in the potential productivity of rangelands in Botswana using the Kalman filter to model productivity change as a stochastic trend. In particular they use the Kalman filter is used to estimate the state of the range as a latent variable in much the same way that the state of technology has been treated as a latent variable by Slade (1989) and Harvey and Marshall (1991). They show that change in the current and long-run equilibrium carrying capacity of an ecological system may be treated as non-stationary trends. This enables them to measure both the speed of return to equilibrium, and the threshold effects that occur when the system loses the capacity to absorb shocks of given magnitude.

probabilistic evolution of the system. This does not stop us from thinking about the effects of changes in institutions or rules. Indeed, it offers a very natural and structured way of doing this. A change in institutions induces a change in policies and hence the probability that the system will develop in particular ways.

We take the policy to be to maximise a measure of social welfare. Specifically, the expected welfare from the policy is:

$$W^u(i) = E^u \Sigma^{\infty}_{t=0} w(x_t, u_t(x_0,\ldots,x_t))$$

The advantage of thinking about the development or evolution of an economy-environment system in this way is that it makes it easy to identify the resilience of the system in a given state. The ways in which decision-makers can influence the process depends on the structure of the system. If the time path for the state variables, x_t, given u_t, is:

$$x_t = f_t(x_0, u_1, \ldots u_t)$$

then the set of all states that are reachable from x_0 at time t depends on the current state of the system and the transition probabilities, P. It is useful to distinguish between transient and recurrent states. States that are revisited are said to be recurrent. Recurrent states are either occupied permanently or revisited periodically, transient states are left after some finite time and never revisited thereafter. It is quite natural to associate recurrent states with the long-term equilibria of a system, and transient states with far-from-equilibrium positions.

To relate this to the concept of resilience, recall that resilience is frequently measured by the size of the disturbance the system can absorb before flipping from one stability domain to another. Now the transition probabilities just described define the probability that a system in one state, and subject to some disturbance regime, will change to another state. This is exactly what the Holling measure requires. However, it is much more general than the Holling measure. It defines the transition probability from one state to another state whether or not that other state lies in a different stability domain.9

The main point here is that the evolutionary potential of an economy-environment system is limited by institutional conditions-or the ‘rules of the game’. The evolutionary possibilities of a system are summarised by the probability law P^u. This depends on institutional conditions, the decision-maker’s objectives, the admissible policies and actions and the strategic behaviour of agents. So for a given set of property rights, a given disturbance regime, and a given state of nature, it may be possible to estimate the

9 Technically, in the special case where P is both irreducible and aperiodic the system will have a unique globally stable equilibrium-it has only one stability domain. In this case the transition probabilities of the system may be said to be equivalent to Holling-resilience measures. In the more general case where P is reducible, the state space may be partitioned into classes corresponding to multiple equilibria. It follows that a sufficient condition for a system to be infinitely resilient is that P is irreducible. If P is reducible (a) the system may, in the limit, occupy any one of a finite number of closed classes; (b) it is sensitive to initial conditions, and (c) it is path dependent (the key properties of complex systems generally). In this case, the limiting transition probabilities of the chain depend on the initial state, i. The future evolution of the system depends on where it starts.

probability that the system will converge by some finite time on some other state of nature. The connection with the notion of sustainability is direct. If the transition probabilities are known, it is possible to estimate either the time the system occupies a particular state (the sustainability of that state) or the time to convergence on any other state (the loss of resilience).

The resilience and hence sustainability of the system in any one state depends on the way it is used-the control policy applied. For most economies the process of development involves a sequence of states. Indeed, early development theory was all about the transition between equilibria—about escaping from states associated with low levels of wellbeing and moving towards states associated with higher levels of well-being. Strategies for sustainability are about enhancing or protecting the resilience of the system in desirable states and reducing the resilience of the system in undesirable states (poverty traps, subsistence or semi-subsistence equilibria and the like) (Perrings, 1998).

4. The dynamics of shallow lakes

Many ecological systems display discontinuities in the equilibria of the state of these systems over time. A famous example is the interactive dynamics of the spruce budworm, its predators and the forest it lives in (Ludwig, Jones and Holling, 1978). As the forest grows, equilibrium budworm numbers are relatively low in the beginning but at a certain point suddenly jump to relatively high numbers. As a consequence, the dynamics of the forest is reversed and living conditions deteriorate, but for a while the budworm density remains relatively high before it returns to low numbers again. This hysteresis effect is due to a non-linearity in the dynamics of the spruce budworm reflecting the role of its predators. It implies that for a range of values for the living conditions, high and low equilibria for the budworm numbers exist with separated domains of attraction. Other examples of ecological systems that display this phenomenon can be found in Ludwig, Walker and Holling (1997).

These hysteresis effects are also important for the interaction between human behaviour and ecological systems. Here we take the example of the eutrophication of shallow lakes (e.g. Carpenter and Cottingham, 1997; Scheffer, 1997) but the situation can be viewed as a metaphor for many of the ecological problems facing us today. Due to increasing agricultural activity, more and more phosphorus is released into the lake. Initially the effect is small but at a certain point in time the lake flips from a so-called oligotrophic state with a relatively high value of ecosystem services to a so-called eutrophic state with a relatively low value. Because of the hysteresis effect, this drop in ecosystem services is only reversible at a high cost because the agricultural activity has to be reduced far below the level where the flip occurred. In other cases the loss of ecosystem services may be irreversible.

Management of the lake has to consider the trade-off between the benefits of agriculture and the benefits of the services that the lake can provide such as fishing, recreation and the use of water for industry and consumption. It will be shown that for a general class of welfare functions optimal management chooses a level of agricultural activity that keeps the lake in an oligotrophic state but can be close to the point where the lake flips to a eutrophic state. This implies that small mistakes can have large costs because the lake can flip and, due to the hysteresis, agricultural activity has to be reduced first below the original level and then increased again to reach the optimal point.

In case such a lake is shared by different communities, each with a welfare function as described above, a game is played between these communities. The services are public and all communities influence the state of the lake by the release of phosphorus from their agricultural activity. It will be shown that typically two Nash equilibria exist. The first one leaves the lake in an oligotrophic state, but it is located somewhat closer to the flip point than the cooperative outcome. The second equilibrium, however, leaves the lake in a eutrophic state with low welfare. If the communities end up in the second Nash equilibrium and decide to coordinate their policies in order to reach the cooperative outcome, the coordination is much more difficult and costly than from the first Nash equilibrium because of the hysteresis. In case the first flip of the lake is not reversible, it is even impossible to reach that point.

Under the assumption that the dynamic processes in the lake are very fast so that the state of the lake adjusts instantaneously to new loading levels of phosphorus, the analysis is essentially static and relatively easy. A full dynamic analysis, however, requires optimal control techniques for non-linear systems. We begin with the steady-state economics of shallow lakes. We then consider the full dynamics. The dynamics prove to be very complex.

The Lake Model

Shallow lakes have been studied intensively over the last two decades and it has been shown that the essential dynamics of the eutrophication process can be modelled by the differential

$$\dot{P}(t) = L(t) - sP(t) + r\frac{P(t)^2}{P(t)^2 + m^2}, P(0) = P_0,$$

equation
where P is the amount of phosphorus in algae, L is the input of phosphorus (the “loading”), s is the rate of loss consisting of sedimentation, outflow and sequestration in other biomass, r is the maximum rate of internal loading and m is the anoxic level (see Carpenter and Cottingham, 1997; Scheffer, 1997). Less is known about deep lakes but from what is known it can be expected that the same model is adequate. Estimates of the parameters of the differential equation for different lakes vary considerably, however, so that a wide range of possible values has to be considered.

By substituting $x = P/m$, $a = L/r$, $b = sm/r$ and by changing the time scale to rt/m, equation

$$\dot{x}(t) = a(t) - bx(t) + \frac{x(t)^2}{x(t)^2 + 1}, x(0) = x_0.$$

(1) can be rewritten as

The last term in the right hand side - the internal loading - has a feature that is the cause of the interesting dynamics. This term is convex over a certain interval and concave over another interval. This means that we loose the convexity of the analysis, the convexity that is at the basis for decentralized decision making. We will see this in more detail in what follows.

In order to understand the essential features of the model suppose that the loading *a* is constant. For high values of *a*, equation (2) always has one stable equilibrium. For lower values of *a*, three things can happen depending on the value of the parameter *b*. If $b \geq 3\sqrt{3}/8$, all values of *a* lead to one stable equilibrium. If $b \leq 0.5$, a range of values of *a* exists, starting at 0, where equation (2) has one high stable equilibrium and one low stable equilibrium, and where the third root of the right-hand side of equation (2) determines the borderline between the two domains of attraction. If $0.5 < b < 3\sqrt{3}/8$, the range with two stable equilibria is preceded by a range of low values of *a* with again only one stable equilibrium. Figures 4-6 show the three different situations.

It is easy to see the hysteresis effect now for $b < 3\sqrt{3}/8$. If the loading *a* is gradually increased, at first the equilibrium levels of phosphorus remain low so that the lake is in an oligotrophic state with a high value of ecological services. At a certain point, however, the lake flips to a eutrophic state with high equilibrium levels of phosphorus and a low value of ecological services. If it is now decided to lower the loading *a* in order to try to restore the lake, it has to be decreased below this flipping point until a point is reached where the lake flips back to an oligotrophic state. If $b \leq 0.5$, the lake is trapped in high equilibrium levels of phophorus, as can easily be seen from figure 5, which means that the first flip is irreversible. In that case only a disturbance of the parameter *b* of the model, e.g. by changing the fauna of the lake, can possibly restore the lake.

Steady-State Economics

Several interest groups operate in relation with the lake modeled in section 2. Since the release of phosphorus into the lake is due to agricultural activity, at least the farmers have an interest in being able to increase the loading. In this way the sector can grow without the necessity to invest in new technology in order to decrease the run off-output ratio. On the other hand, a clean lake is preferred by fishermen, drinking water companies, other industry that makes use of the water, and people who spent their leisure time on or along the lake.

There are many potential conflicts that can arise on the use of the lake. The conflict that one perhaps first of all would think of is the conflict between farmers (and in extension consumers of agricultural products) and those who demand an oligotrophic lake in order to sustain their well-being. We will come back to this conflict later. However, for now we will instead focus on conflicts between different communities.

Suppose that a community or country, balancing the different interests, can agree on a welfare function of the form $\ln a - cx^2$. The lake has value as a waste sink for agriculture and it provides ecological services that decrease with the amount of phosphorus in algae. Furthermore, suppose that the lake is shared by *n* communities or countries with the same welfare function. If the amount of phosphorus adjusts instantaneously to the loading levels chosen by these communities without costs, just a static steady-state problem results. Suppose also that $b = 0.6$ in equation (2), so that according to the above analysis, the lake displays hysteresis but a flip to a eutrophic state is reversible. This is the most interesting case but the analysis for other values of the parameter b can of course be performed in a similar way.

First we will concentrate on the optimal management problem which implies the maximization of $S_i \ln a_i - ncx^2, c > 0,$ subject to

$$\sum_{i=1}^{n} a_i - 0.6x + \frac{x^2}{x^2+1} = 0.$$

The logarithmic functional form will lead to an outcome that is independent of n, so that it can be a benchmark for Nash equilibria regardless the number of communities. It is also assumed that the area around the lake is large enough, so that adding new communities does not lead to crowding out and the objectives can be additive in the number n.
Simple calculus shows that the optimal steady-state amount of phosphorus x^* is determined

$$0.6 - \frac{2x}{(x^2+1)^2} - 2cx(0.6x - \frac{x^2}{x^2+1}) = 0.$$

by
If $c = 1$, this yields $x^* = 0.33$, with the optimal steady-state loading $\Sigma a_i^* = 0.1$. It means that the lake is in an oligotrophic state but note that it is also possible to end up in the eutrophic state $x = 1$ for the same level of loading, which will happen if the initial amount of phosphorus is in the upper domain of attraction.

Optimal management, of course, does not necessarily have to lead to an oligotrophic state of the lake. If the welfare function attaches a relatively low weight to ecological services it can become optimal to choose a eutrophic state with a high level of agricultural activities. In the steady-state objective, c denotes the relative weight of the loss of ecological services with respect to the value of the lake as a waste sink for phosphorus. For large values of c, the optimal steady-state problem has one maximum for an x below the flipping point. As c decreases, first a local maximum appears for a high x while the global maximum is still reached for a low x, but for c low enough ($c \leq 0.36$) the global maximum occurs for a high x beyond the flipping point. In the sequel of this paper it is assumed that enough weight is attached to the services of the lake, so that it will be optimal to aim for an oligotrophic state. This is guaranteed by taking $c = 1$.

Flipping occurs when total loading is increased to $a = 0.1021$, which means that the lake will be managed not far from what can be called the "edge of hysteresis" (Brock, Carpenter and Ludwig, 1997). Small perturbations that cause a flip will have high costs, not only directly by a jump to a high x but also indirectly because of the long path of return to the optimal situation. Therefore, policy considerations might lead to a precautionary principle.

Suppose now that the communities or countries do not cooperate in optimally managing the lake so that it is appropriate to search for Nash equilibria for the objectives $\ln a_i - cx^2$, $i = 1,2,...,n$, subject to (3). Simple calculus shows that the steady-state amount of phosphorus in a

$$0.6 - \frac{2x}{(x^2+1)^2} - \frac{1}{n} 2cx(0.6x - \frac{x^2}{x^2+1}) = 0.$$

Nash equilibrium has to satisfy
For the sake of exposition we take n = 2 and c = 1, so that optimal management of the lake leads to the oligotrophic state x^* = 0.33 with optimal loading levels a_i^* = 0.05, i = 1,2 (see above).

Equation (5) has three solutions but only two of them relate to Nash equilibria. The first one is x = 0.36 with equilibrium loading levels a_i = 0.0506, i = 1,2. This equilibrium point lies between the full cooperative outcome and the flipping point which shows that non-cooperative behaviour not only leads to a standard loss in welfare but also brings the lake closer to the edge of hysteresis.

More interesting is that also another Nash equilibrium exists that yields the steady-state amount of phosphorus x = 1.51 with equilibrium loading levels a_i = 0.1054, i = 1,2. If the communities focus on this equilibrium point, the lake is in a eutrophic state. Welfare is much lower here. In the full cooperative outcome the communities each have a welfare level of -3.1068, and in the Nash equilibrium with an oligotrophic state the welfare level of each community is -3.1134, but in the Nash equilibrium with a eutrophic state the welfare level is -4.5301.

Each community will have some policy, like a tax on emissions, in order to regulate the release of phosphorus from agricultural activities to a desired level. If the communities are in the first Nash equilibrium and decide to cooperate, it is relatively easy to redesign the policy in order to regulate the loading to the optimal level. However, after a flip of the lake has occurred and the communities are locked into the bad Nash equilibrium, it is much more difficult for the people to reach the cooperative outcome, in reaction to a higher tax. It is not enough to reduce the loading levels from a_i = 0.1054 to a_i^* = 0.05, $i = 1,2$. The reason is that the lake will then still be in a eutrophic state with the steady-state amount of phosphorus x = 1, because the adjustment process started in the upper domain of attraction. What is needed is first a further reduction to the loading levels a_i = 0.0449, $i = 1,2$, where the lake flips back to an oligotrophic state. After this flip back has occurred, the loading levels can be increased again to a_i^* = 0.05, $i = 1,2$. The two stages in the reaction process to higher taxes are a consequence of the hysteresis in ecological systems. Lack of care in the neighbourhood of the flipping point may lead to the bad Nash equilibrium and may, therefore, considerably increase the costs of restoring cooperation.[10]

Dynamic analysis

The problem we wish to consider here is the existence of shadow prices for lake problem. Flips from one basin of stability to another take time (one to three years for some lakes). There are accordingly no discontinuities, but the problem is a dynamic one. The static

[10]Another way of saying the same thing is that transaction costs increase because of the existence of non-convexities.

analysis of the previous section is not sufficient. In order to find the optimal management strategy, we need to take the transients into account. The way to do this is to find the time path of the run off that will maximize the present value of future utility. This can be analysed by using Pontryagin's maximum principle (Brock and Starrett. 1999; Mäler, Xepapadeas, and de Zeeuw (1999)[11].

Remember that the maximum principle introduces an auxiliary or co-state variable that can be interpreted as a shadow or accounting price on the state variable. In the lake case, there will be a price on the stock of phosphorous in the algae in the lake. This price will change as follows

$$\frac{dp}{dt} = [\ d + b - \frac{2x}{(1+x^2)^2}]p$$

where δ is the utility discount rate.

$$1/a + p = 0.$$

Thus the price must be negative, reflecting the fact that *x*, the amount of phosphorous in the lake is a bad. By using this relation we can rewrite the equation for the dynamics of "p" as an

$$\frac{da}{dt} = -[(b + d) - \frac{2x}{(1+x^2)^2}]a + 2cxa^2$$

equation of the dynamics of "a":
By combining the equation da/dt = 0 and the equation for dx/dt = 0, we are able to characterise the optimal management of the phosphorous flow into the lake.

That there are three points of intersections between the curves is arbitrary and depends on the parameters. There can be one, three, or even a countable number of equilibria, but the generic case will have a finite odd number of equilibria under plausible economic restrictions. These intersections between the two curves indicate potential equilibria (see figure 6). However, the middle one is unstable and can be ruled out. The remaining two are saddlepoints and may be steady states. To which will the optimum path converge? That depends on the initial stock of phosphorous in the lake. If the lake is much polluted initially, the eutrophic steady state will be the ultimate outcome, while if the initial pollution is low, the optimal path will bring the lake to an oligotrophic equilibrium. Which equilibrium the optimal path will converge on is determined by the initial stock of phosphorous in the lake. If that level exceeds a certain level x_S (the so called Skiba point after the mathematician who first studied these convex-concave dynamic problems), the system will converge to the eutrophic equilibrium and vice versa.

[11] Mäler, Xepapadeas and de Zeeuw, (1999) offer a full simulation of the lake model with both open loop and feed back equilibria for a differential game on the use of the lake.

Assume now that phosphorous flows into the lake. We may now prepare a phase diagram of the system by plotting "*a*" on the vertical axis and "*x*" on the horizontal. The curves *da/dt*=0 and *dx/dt*=0 divide the phase diagram into four of regions: (1) dx/dt<0,da/dt<0; (2) *dx/dt*<0,*da/dt*>0; (3) *dx/dt*>0,*da/dt*<0; (4) *dx/dt*>0,*da/dt*>0. By visual analysis of these four regions and the directions of movement of the variables "*a*" and "*x*" in each, it is possible to locate all initial pairs (x,a) such that the dynamics converges to a steady state in both "forwards" and "backwards" time. Candidate optimal paths are located on the phase diagram as those which converge to a steady state in forward time. There are techniques available Skiba (1978), Dechert and Nishimura (1983), Brock and Malliaris (1989) that one can use to compare the value of the objective on each of these candidates without having to do the actual integration required to explicitly evaluate the objective.

Using this type of analysis, Brock and Starrett (1999) provide a complete analysis of the location of the optimum for the case-call it "1"-where there is only one steady state equilibrium and for the case-call it "2"-where there are only three steady state equilibria. They indicate how to generalize this analysis to the generic case of any odd number of equilibria. In case 1, where there is only one steady state equilibrium, the analysis is standard. One locates the optimum trajectory on the phase diagram by locating for each initial x, the value a(x) such that the pair of differential equations converges to the steady state as time advances.

Consider case 2 where there are three equilibria. The low x and high x equilibria are saddle points, whereas the middle x equilibrium is an unstable spiral. The middle steady state can always be shown not to be an optimum. Notice that in Case 2 there are always initial values of x where there are two candidate optima. For each initial x there are two values of a, one low, one high, such that the trajectory starting from each of these converges to one of the steady states. More particularly, a high a leads to high steady state x and low a leads to low steady state x. Case 2 divides into three subcases: 2.1. It is always optimal to go to the low x steady state, no matter how big the initial x is. 2.2. It is always optimal to go to the high x steady state no matter how big the initial x is.2.3. There is a cutoff point, the "Skiba" point, such that if initial x is below (above) the Skiba point, the optimal path converges to the low x steady state (above). The paths yield the same value at the Skiba point.

Let us use the above discussion of our very stylized model to inquire into changes that have to be made in conventional "convex" economics to deal with this type of situation where hysteresis and irreversibility may be endemic due to non-convexity of the state dynamics. We do this in several parts. First we consider the case where the utility of loading u(a) (which we are setting equal to log(a) for illustrative purposes) is generated by the sum of profits of firms located in the watershed of the lake and where the farmers are numerous enough that strategic interactions may be ignored. We consider the workability of taxation of loadings at marginal social cost.

Parenthetically we remark that in practice taxes may not be the most efficient instrument, especially if restoration of the lake is more easily done by ecological sequestration of harmful materials than by stopping pollutants from getting into the lake in the first place. Sequestration might be done more efficiently by, for example, co-payment schemes for buffer banks along stream course ways and riparian reserves rather than direct taxation of loadings. This is likely to be the case when (i) sequestration is quite easily done by inducing direct effort to sequester runoff nutrients by devices such as buffer banking streams, (ii) the elasticity of substitution between taxed inputs and other inputs is low in production functions, and (iii) risk aversion of individual firms is high (because of incompletenss of hedging markets and the tax translating into a fixed cost equivalent via low substitution elasticity and high administrative costs, perhaps).

It is beyond the scope of this chapter to analyze other instruments of regulation other than to stress to the reader that other instruments should be considered in any practical application. Return now to discussion of implementation of taxation on loadings. Our initial remarks at the beginning of this chapter stressed the role of convexity of the abstract production technology set in the design of conventional decentralized regulatory schemes such as taxation at marginal social cost at each point in time. We discuss here the workability of this type of decentralization in the above setting where not only are there two cases but case 2 has three subcases. Note that except for case 2.3 with initial x at the Skiba point, there is a unique socially optimum path, call it $\{a^*(t), x^*(t), x^*(0)=xo$ given$\}$. Let each firm f in the lake's watershed have profit function $u(a(f),f)$ where firm f loads $a(f)$ into the lake. Let $u(a)$ denote the maximum of the sum of $u(a(f),f)$ over all firms f, subject to the constraint that the $a(f)$ sum to less than or equal to a. This is a concave problem provided each u is concave increasing in $a(f)$. Hence, under regularity conditions on each function $u(.,f)$ we can implement decentralization schemes using conventional economics for each target level of total loading. The regularity condition needed is that the derivative of each u be decreasing in a, be very large for very small a, and very small for very large a. That is, the marginal product of loading must fall with loading and be very large (very small) for very small (very large) loadings.

Notice that a linear utility has constant marginal product and, hence, will not satisfy the regularity condition. But $u(a)=log(a)$ satisfies it. Under the regularity assumption-call it the “controllability condition”-we can induce each firm to load an amount $a^*(f,t)$ that sums across firms to the socially desired target total loading $a^*(t)$. This may be done by imposing a tax $T^*(t)$ on each unit of loading at marginal social cost $-p^*(t)$ where $p^*(t)$ is the shadow price of x evaluated along the socially optimal path $(a^*(t), x^*(t), x^*(0)=xo$ given). p*(t) is negative because it is the derivative of the optimal value function w.r.t x and x is a bad. Indeed if each $u(.,f)$ is controllable we can induce any choice of $a(f)$ we wish by charging an appropriate "tax-price" T. The tax $T^*(t)$ can be written as a continuous function of the state variable $x^*(t)$ for case 1 and cases 2.1, 2.3 but not for case 2.2. But this discontinuity in the tax function presents no problem for the optimum planner in decentralizing this community of firms. All they need to do is to design incentives to induce them to produce the total target socially optimal loading $a^*(t)$ at each point of time t. The resulting system solves the differential equation

$$dx/dt=a^*(t)+g(x$$

$x(0)=xo$ given. But x^* itself solves the differential equation

$$dx^*/dt=a^*(t)+g(x^*)$$

x(0)=xo* given. A trivial adaptation of the usual argument for uniqueness of solutions to differential equations proves that solution *x* equals solution *x** since the control *a** is the same. Even the case 2.2 can be decentralized. For initial *x*'s not equal to the Skiba point, simply follow the above. For initial *x* equal to the Skiba point simply make up one's mind which path (they both have the same welfare) you want to follow and design *T*(t)* as above to induce *a*(t)* for the total loading along that chosen path. This treatment looks more special than it really is. For example imagine each firm has a vector of state variables (i.e. slow variables) which are costly to rapidly adjust as well as flow variables like the above. As long as each firm's problem is strictly concave we may form a global overall optimization problem, optimize it to get the optimal loading of the lake, tax loading of each firm at marginal social cost *T*(t),* present each firm with this tax schedule and tell each to go ahead and maximize the capitalized value of their profit stream net of taxes. Since their optimum problem is still concave, they will reproduce their part of the socially optimum path of loading as above.

It follows that introduction of slow and fast variables on the firm side presents no problem for standard decentralization theory so long as the firm's problem remains concave in the relevant arguments and an analogue of the controllability condition holds. Spatial scale considerations present no problem in this idealized full information deterministic world either. If there are multiple watersheds, the overall social optimization problem may be solved, and taxation may be imposed on the firms at marginal social cost as before. As before, so long as the firms' problems are concave in their state and control variables, and a dynamic analogue of the controllability condition is satisfied, the tax schedule may be designed so that firms are induced to follow the overall social optimum.

The introduction of uncertainty presents no problem provided there is common agreement amongst the firms and the regulator on the true distribution of stochastic shocks to the system and provided that the firms face concave optimization problems and a stochastic analog of the controllability condition is satisfied. Hence, the essence of the tax problem from an abstract theoretical point of view is that the targets of the taxation, i.e. the firms, all face concave optimization problems where an analogue of the controllability condition holds. The overall social optimization problem of the coupled economic ecological system can be non-concave (non-convex overall production set in the language of welfare economics exposited at the beginning of this chapter). What is key to successful decentralization of incentives is whether the system may be decomposed so that the targets of the decentralized regulatory scheme themselves face concave problems which are "controllable" (the firms in our case).

In this case the targets can be "controlled" via dynamic state dependent tax schedules to reproduce the socially optimal values of their inputs into the overall economic/ecological system. The ecological system, even though its dynamics are non-convex, will reproduce the socially optimal path for the state variables, given that the inputs from the economic side are controlled at their socially optimal levels. Although what we have said above is sound from the point of view of pure theory it needs revision and supplementation for the complexities of actual practice. In actual practice, the dynamics of both the economic system and ecological system are not known and must be estimated. The distribution of stochastic shocks must also be estimated. Ludwig (1994) studies a harvesting problem and treats several levels of uncertainty including difficulties in measuring the state, uncovering the true dynamics, administering control, and uncovering the true distribution of outside shocks to the system. He also treats not only the case of continually occuring small shocks but also the case of rare but large shocks. He argues that precautionary principles tend to get strengthened as one adds more layers of uncertainty.

The optimal design of regulatory instruments becomes much more complicated in this more realistic world. However if there is a presumption that the economic side of the coupled economic/ecological system displays enough "regularity" in the sense of concavity of production, (i.e. convexity of production sets), and quasi-concavity of utility functions, then we may borrow from a large literature in economics to induce agents to choose in a decentralized manner whatever loading choice is deemed optimal from the ecological side. The extra complications created by scientific ignorance about the ecological dynamics, stochasticity, ecological non-convexities, cross scale interactions, slow and fast variables, unobserved slow variables etc. may manifest themselves much more in the decision about the amount of loading to be tolerated than in the exact manner how that chosen amount of loading is to be implemented by the agents of the economy. Recent work on robust control may be a useful way to extend the work discussed here. Robust control theory gives a precise way of modeling the type of risk present when mis-specification error looms large and the social planner wishes to insure against mis-specification risk when designing regulatory policy.

Extension and application of this literature to misspecification of non-convex ecological dynamics and the impact of this misspecification of non-convex dynamics with potential alternative stable state upon regulatory design seems to us to be a very worthwhile research project. Return now to discussion of optimal taxes when model specification is correct. Notice that the optimal tax is the marginal social cost of loading which depends upon the state (i.e. the fully specified history in the stochastic case) of the system. In our simple time stationary recursive system, this tax can be specified as a time stationary function of the state "x" of the system, but in general non-stationary stochastic settings the state description will be more complicated. Complexities of practical implementation of such detailed contingent specification of policies argue for implementation of simple approximate policies. This consideration argues that computational work like Dechert (1999) which computes social welfare from simple policies and evaluates the cost of simplicity, will be valuable.

Economists commonly argue in favor of tax instruments over many other modes of regulation such as quotas because of flexibility. However, in practice there are many complications that may modify this prescription: (i) Firms may face overhead costs which must be covered or they will go out of business. This generates a non- convexity. (ii) Firms may have dimensions of adjustment to incentives that are not captured by estimates of their technology used in setting the level of taxes, quotas, loading reduction co-payments, buffer bank co-payments, or any other regulatory instrument. (iii) Any mode of regulation induces costs on the regulated as well as the regulator. These administrative costs as well as heterogeneity of firm types argue for regulatory tiering in practice (Brock and Evans (1986)).

Furthermore, it is common to find that a small number of firms to cause the bulk of the problem. This is suggested by Gibrat's Law-that the size distribution of firms in actual practice is roughly lognormal.12 Hence, one might avoid both political and administrative costs by regulating only the small number of operations that are responsible for the bulk of the problem, and simply leave the rest of the firms alone.

The small numbers case

We turn now to the case where there is a small number of strategically interacting firms in the lake's watershed. We sketched a theory above that suggests that regulation of a large number of firms (a large enough number so that incentives to act strategically are minimal) might be easily decentralized because each firm is solving a concave, tax-controllable, problem at each point in time. Hence all the decentralizer has to do is design a system where the firm sector is induced to load the target $a^*(t)$ at each point in time t. This can be done with taxes, but could also be done with emission permit markets where a firm must have a permit to emit at any date t, a permit to emit lasts only one "period", the permit market is open each "period", the agency sets $a^*(t)$ permits to be sold each "period." Of course this is very idealized but it suggests that decentralized regulation may be achievable for large numbers cases where strategic interactions may be ignored. Decentralization is more difficult in the small numbers case discussed above in this section because each operator took into account the impact of the others' actions on the dynamical state equation.

We sketched the case of an open loop Nash equilibrium above. We turn to a very brief sketch when it is appropriate to focus on other concepts of equilibrium. First there is the "Ostrom" issue (Ostrom,1990; Ostrom et al, 1994) of locating conditions where the small group might self-evolve institutions that do better than Nash non-cooperative equilibrium. For example the temporal and spatial scale of the problem may interact with the determinants of the ease or difficulty of organizing collective action as detailed by Olson's classic on collective action (Olson, 1965) and Ostrom's work on Common Property Resource (CPR) management. For example Ostrom (1990) lists factors (and cites case studies to back them up) that are positively associated with success of groups at self organizing the provision of public goods, self organizing avoidance of tragedies of the commons, self managing Common Property Resources (CPR's), and the like. Hence, the use of coercive schemes or other governmental catalyzed schemes are not required to get a workable solution.

Ostrom's analysis is germane to our case of locating features of the underlying social, cultural, and ecological context that might allow our community of non-cooperative Nash players to do better. Of course the literature on supergames and dynamic games suggests that small discounting of the future utilities allows threat strategies to be designed that will support a self-enforcing agreement to manage the system at the optimal level. Intuitively the non-cooperative Nash equilibrium is returned to and played forever

12 See Evans's work in Brock and Evans (1986) for studies of Gibrat's Law and studies of regulation of pollution and other negative externalities in actual practice.

if anyone is caught loading more than the globally optimal loading. It is beyond the scope of this expository piece to get into the issues raised by the literature on dynamic commons games (see, for example, Dutta and Radner (1999) for most recent work). We turn to the less formal approach of Ostrom (1990), Ostrom et al. (1994). Here is a list of Ostrom's (1990) conditions that are positively associated with successful self-evolution of institutions for managing a common-property resource like a lake which give social payoff better than non-cooperative equilibrium:

1. Most appropriators share a common judgment that they will be harmed if they do not adopt an alternative rule.
2. Most appropriators will be affected in similar ways by the proposed rule changes.
3. Most appropriators highly value the continuation activities from this CPR; in other words, they have low discount rates.
4. Appropriators face relatively low information, transformation, and enforcement costs.

5. Most appropriators share generalized norms of reciprocity and trust that can be used as initial social capital.
6. The group appropriating from the CPR is relatively small and stable." (Ostrom,1990, p. 211).

The word appropriator refers to a user of a CPR, but could equally apply to self-organized communal users of a commonly shared resource like the watershed of a lake and the lake itself. In this case, enforcement costs refer to detection and policing of “shirkers”. Ostrom lists many case studies to document the importance of parameters 1-6 which she has numbered in order of importance with 1 being the most important for successful self-organization of “workable” cooperation. Second, we used the concept of open loop dynamic Nash non-cooperative equilibrium above. Dechert (1978), (1999) has produced a handy method for computing open loop Nash equilibria that applies to our setting. He produces an optimal control problem whose solution is a Nash equilibrium. For our case, his "as-if" control problem is the sum of the loading utilities across the players minus the cost for one player (recall that the cost is the same for all players). Hence, (generically), we have an odd number of steady state equilibria and in the case of one or three steady state equilibria the above taxonomy of case 1 and case 2 with three sub-cases applies.

In Dechert and Brock (1999) he uses the “upwind Gauss-Seidel” method to devise a very rapid algorithm to compute solutions to a discrete time version of our lake problem.13 For $u(a)=log(a)$, $c(x)=x^2$, $dx/dt=a-bx+g(x)$, $g(x)=x^2/[(1+x^2)]$, and parameter values roughly the same as the lake example we discussed above, he shows the following. First, for two player games in Nash non-cooperative open-loop equilibrium, overloading relative to the social optimum loading at each level of x is surprisingly small until initial x is past the point of inflection of $g(x)$. This is the level of x where positive feedback effects are triggered in the ecosystem dynamics. However, for three or more players there is an abrupt change in the level of shirking even before the inflection point of $g(x)$ is reached.

13 Up-wind Gauss-Seidel may be used because the problem has a one dimensional state variable, is recursive, and has a monotone optimal policy function. These features may be exploited to yield a very fast computational algorithm.

Second, Dechert sets a constant tax at the “level necessary to make the social optimal steady state the steady state solution to the dynamic game” (Dechert and Brock (1999, p. 12)). He shows numerically that the dynamic game solution converges to the social optimum steady state even up to 100 players.

6 Concluding remarks

The shallow lake case shows evidence of hysteresis in the state dynamics, implying that a system that flips from one state to another at some value of the control may require a very different value of the control to return it to the original state. Because such systems can flip suddenly from one state to another as a result of particular events ‘events’, it is difficult for any regulatory agency to observe the signals of an impending change in time to take action to avert it. In a managed system, the dynamics of the resource are revealed through the response of the state variables to the controls. The closer the system is brought to the boundaries of the stability domain, the higher is the risk of an unanticipated irreversible or only slowly reversible change as the system flips from a higher productivity state to a lower productivity state. The same phenomenon makes it difficult to devise a decentralised regulatory system involving taxes or charges. Once the system has flipped to an undesirable state, taxes or user fees would have to be such as to drive runoff well below the original levels and to hold them there in order to overcome the hysteresis effect.

Systems ranging from coral reefs to semi-arid savannas have been observed to behave in very similar ways. From a regulatory perspective the problem is precisely that bifurcation points may not been seen before they are reached. The observable level of environmental quality does not generally offer a reliable indicator of the system's relative position with respect to thresholds. Moreover, the conditions under which ecosystems respond to increasing stress without suffering an irreversible or near irreversible loss appear to be fairly restrictive.

Finally, it is worth noting that modelling the interactive decisions of human agents and other species poses special problems. Human decision-makers do have the capacity to be forward looking. In this they seem to differ from other species. More importantly, they are capable of social learning and possess institutional memory (libraries etc). This is obvious but it is hard to write down differential equations in ways that differentiate human interactions from the interactions in other species in terms of differential equations. In economics, we typically capture that difference in the level of sophistication agents are assumed to have in their expectations of the future and their strategic behaviour. For example in rational expectations models we deal with systems of ordinary differential equations where we solve the system "forwards" whereas in other sciences, including mathematical biology, we deal with systems where we solve the ordinary differential equations "backwards." It may be argued that the economic approach is a benchmark way of incorporating the relative ingenuity of humans. The appropriate level of complexity of forward looking behaviour and level of strategic interactions amongst human beings relative to animal and plant community systems demands a correspondingly more sophisticated modelling approach than for example game theoretic concepts such as ESS used by J. Maynard Smith. Economic applications of game theory for example, exploit more levels of iterated common knowledge than do the biological and evolutionary applications. The biological and evolutionary applications also do not exploit the degree of farsightedness as does the rational expectations literature. It can be argued that this is entirely appropriate given the level of sophisticated cognition contained in human beings relative to other species.

References

Brock W.A. 1998. Ideas on quantification of resilience-based management, Environment and Development Economics 3(2): 239-243.

Brock, W. and Evans, D., 1986, The Economics of Small Businesses: Their Role And Regulation in the U.S. Economy, New York: Holmes And Meier.

Brock, W.A., and A.G. Malliaris, 1989, Differential Equations, Stability and Chaos in Dynamic Economics, North Holland, Amsterdam.

Brock, W.A., and D. Starrett, 1998, "Nonconvexities in ecological management problems," Mimeo.

Brown, D., 1991, Equilibrium Analysis with Non-Convex Technologies, In Hildenbrand, W., Sonnenschein, H., Eds., Handbook Of Mathematical Economics: Volume IV, North Holland: Amsterdam.

Carpenter S. and Pace M., 1997. Dystrophy and eutrophy in lake ecosystems: implications of fluctuating inputs, Oikos 78: 3-14.

Carpenter S. Ludwig D. and Brock W.A. 1999. Management of eutrophication in lakes subject to potentially irreversible change, Ecological Applications, in press.

Carpenter, S., Brock, W., Hanson, P., 1999, Ecological a nd Social Dynamics in Simple Models of Ecosystem Management, Conservation Ecology 3(2).

Carpenter, S.R., and K.L. Cottingham, 1997, Resilience and restoration of lakes, Conservation Ecology 1:2.

Coase, R. 1960, The problem of social cost, Journal of Law and Economics 3: 1-44.

Common M. and Perrings C. 1992. Towards an ecological economics of sustainability, Ecological Economics 6: 7-34.

Conway G. R. 1993, Sustainable agriculture: the trade-offs with productivity, stability and equitability, in Barbier E. B. (ed.) Economics and Ecology : New Frontiers and Sustainable Development, Chapman and Hall, London.

Dasgupta P. and Mäler K.-G. 1995. Poverty, institutions and the environmental resource base. In Behrman J. and Strinivasan T.N. eds. Handbook of development economics 3(A). Elsevier, Amsterdam.

Dasgupta, P. and K.-G. Mäler, 2000, Net National Product as a Measure of Social Well-Being, Environment and Development, in press.

Dechert, W. and Brock W.A. 1999. details to come.

Dechert, W. and Nishimura, K., 1983, A Complete Characterization of Optimal Growth Paths in an Aggregated Model With a Non-Concave Production Function, Journal Of Economic Theory, 19, 332-354.

Dechert, W., 1978, Optimal Control Problems From Second Order Difference Equations, Journal of Economic Theory, 19, 50-63.

Dechert, W., 1999, Non Cooperative Dynamic Games: A Control Theoretic Approach, Department of Economics, University of Houston.

Dutta, P. and Radner, R., 1999, Work In Progress On Global Warming Games, Department of Economics, Columbia University; Stern School Of Business, New York University.

Elton C.S. 1975,. The ecology of invasions by animals and plants, London, Chapman and Hall.

Holling C. S., D. W. Schindler, B. W. Walker, and J. Roughgarden (1995) Biodiversity in the functioning of ecosystems : an ecological primer and synthesis. In Perrings C. A. , K.- G. Mäler, C. Folke, C. S. Holling, and B.- O. Jansson (eds.), Biodiversity Loss: Ecological and Economic Issues, Cambridge University Press, Cambridge: 44-83.

Holling C.S. 1973 Resilience and Stability of Ecological Systems. Annual Review of Ecological Systems 4: 1-24.

Holling C.S. 1986 The Resilience of Terrestrial Ecosystems: Local Surprise and Global Change. In W.C. Clark and R.E. Munn, eds. Sustainable Development of the Biosphere. Cambridge, Cambridge University Press.

Levin S.A., Barrett S., Aniyar S., Baumol W., Bliss C., Bolin B., Dasgupta P., Ehrlich P., Folfe C., Gren I.-M., Holling C.S., Jansson A.-M., Jansson B.-O., Martin D., Mäler K.-G., Perrings C. and Sheshinsky E. 1998. Resilience in natural and socioeconomic systems, Environment and Development Economics 3(2):222-234.

Ludwig D., Walker B.H. and Holling C.S. and 1997. Sustainability, stability and resilience, Conservation Ecology 1(1): 7

Ludwig, D., 1994, A Theory Of Sustainable Harvesting, Siam Journal of Applied Mathematics, 55: 564-575.

Ludwig, D., D.D. Jones and C.S. Holling, 1978, Qualitative analysis of insect outbreak systems: the spruce budworm and forest, Journal of Animal Ecology 47, 315-332.

Mäler K.-G., Xepapadeas A. and de Zeeuw A. 2000. The Economics of Shallow Lakes, mimeo.

May R.M. 1973. Stability and complexity in model ecosystems, Princeton University Press, Princeton.

McNaughton S.J. 1993. In Schulze E.D. and Mooney H.A. eds. Biodiversity and ecosystem function, Springer, Berlin: 361-384.

Olson, M., 1965, The Logic of Collective Action: Public Goods and the Theory of Groups, Cambridge, Mass.:Harvard University Press.

Ostrom, E., 1990, Governing The Commons: The Evolution of Institutions for Collective Action, New York: Cambridge University Press.

Ostrom, E., Gardner, R., Walker, J., 1994, Rules, Games, and Common-Pool Resources, Ann Arbor: University Of Michigan Press.

Pearce, D. and G. Atkinson 1993, Capital Theory and the Measurement of Sustainable Development: An Indicator of Weak Sustainability, Ecological Economics, 8(2): 103-108.

Perrings C. 1995. Ecological Resilience in the Sustainability of Economic Development, Economie Appliquée 48, 2: 121-142.

Perrings C. 1998. Resilience in the dynamics of economy-environment systems, Environmental and Resource Economics 11(3-4): 503-520.

Perrings C. and Walker B.H. 1997. Biodiversity, resilience and the control of ecological-economic systems: the case of fire-driven rangelands, Ecological Economics 22(1):73-83.

Pezzey J.C.V. 1997. Sustainability constraints versus ‘optimality’ versus intertemporal concern, and axioms versus data, Land Economics 73(4): 448-466.
Pimm S.L. 1984 The complexity and stability of ecosystems. Nature 307: 321-326.
Reed W.J. 1979. Optimal escapement levels in stochastic and deterministic harvesting models. Journal of Environmental Economics and Management 6: 350-363.
Reed W.J.(1988) Optimal harvesting of a fishery subject to random catastrophic collapse. IMA Journal of Mathematics Applied in Medicine and Biology 5: 215-235.
Scheffer, M., 1997, Ecology of Shallow Lakes, Chapman and Hall, New York.
Skiba, A.K., 1978, Optimal growth with a convex-concave production function, Econometrica 46, 3, 527-539.
Tilman D., Wedin D. and Knops J. 1996. Productivity and sustainability influenced by biodiversity in grassland ecosystems, Nature 379: 718-720.
Tsur Y. and Zemel A. 1994. Endangered species and natural resource exploitation: extinction versus coexistence, Natural Resource Modeling 8(4): 389-413.
Tsur Y. and Zemel A. 1997. On resource management under uncertainty: the case of pollution control. Paper presented at the 8th EAERE Conference, Tilburg,.
Westoby M. Walker B. and Noy-Meir I. 1989. Opportunistic management for rangelands not at equilibrium. Journal of Range Management 42(4): 266-274.
World Bank 1996, Expanding the Measure of Wealth: Indicators of Environmentally Sustainable Development.Washington, D.C.: World Bank.