



TI 2009-033/1

Tinbergen Institute Discussion Paper

Shallow Lake Economics Run Deep: Nonlinear Aspects of an Economic- Ecological Interest Conflict

Florian Wagener

CeNDEF, Dept. of Quantitative Economics, Faculty of Economics & Business, University of Amsterdam, and Tinbergen Institute.

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Shallow lake economics run deep:

Nonlinear aspects of an economic-ecological interest conflict

Florian Wagener¹

1: CeNDEF, Dept of Quantitative Economics, University of Amsterdam

Address

Florian Wagener, CeNDEF, University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands; phone: +31 20 525 4034; fax: +31 525 4349; email: F.O.O.Wagener@uva.nl

Keywords

Shallow lake, optimal management, dynamic games, bifurcation analysis

Abstract

Outcomes of the shallow lake interest conflict are presented in a number of different contexts: quasi-static and dynamic social planning, and quasi-static one-shot and repeated non-cooperative play. As the underlying dynamics are non-convex, the analysis uses geometrical-numerical methods: the possible kinds of solutions are efficiently classified in bifurcation diagrams.

1. Introduction

In 2003, Mäler, Xepapadeas and de Zeeuw analysed the economics of the pollution of shallow lakes^{1,2,3,4}. They did this in two contexts: the first is optimal management problem, where a single social planner has to weigh the benefits of agricultural activities, which indirectly pollute the lake, against the costs arising from decreasing ecosystem services of the lake. In the second context, several social planners - communities, countries - share access to a lake, and they play some kind of pollution game against each other. They moreover investigated how the outcome of the game can be improved by setting an appropriate tax scheme.

Very quickly, it became apparent that the shallow lake problem class has a rich structure. Usually in economic analysis situations are considered where the dynamics of the stock variable is assumed to be a convex function of the state of the system, and a linear function of the action of the agents. In that case, if the agent chooses a time-constant action, there is a unique steady-state value to which the stock can converge.

It has been recognised already long ago that in many ecological systems the situation is different, as there may exist tipping points⁵; these are critical values for the action, at which the ecosystem changes, or 'flips', in a way that is not instantly reversible. The system is said to undergo a catastrophic shift. Such shifts have been detected and documented in a great variety of systems⁶; the flipping mechanism of shallow lakes is detailed below⁷. One of the consequences of the existence of flipping points is that for a given action, there may be multiple steady state values of the stock, with radically different characteristics.

Whenever human activity influences the state of an ecosystem, usually through some form of pollution that is a by-product of some kind of production activity, the difficult problem arises of assessing the relative interests of producers affecting the ecosystem and producers and consumers enjoying it. To obtain a benchmark, which should be seen as the best possible scenario, the existence of a benevolent and omniscient social planner is postulated that tries to maximise society's welfare, taking into account the interests of both groups. This, of course, would be an entirely classical economic problem, if the dynamics of the pollution stock would be convex. It being non-convex changes matters considerably.

Standard economic problems usually feature a single equilibrium, and most of the results are about the *degree* of change of that equilibrium if certain parameters are varied. In a problem featuring multiple possible steady states, the *kind* of solution may change if parameters are varied: actually there can be huge qualitative changes in the outcome. Consequently, it is necessary to classify the outcomes of the economic analysis by their qualitative characteristics: are the interests of the producers to be preferred, and is the ecosystem to function mainly as a waste dump, or is it better for society to restrict production activities and enjoy the ecosystem. For this kind of analysis, bifurcation theory is an appropriate tool, as it provides the mathematics of qualitative change. A bifurcation analysis results in a bifurcation diagram, which classifies for different values of the parameters the qualitative features of the outcome.

In the following, some recent work on the economics of shallow lakes will be reviewed, presenting the results as much as possible in bifurcational form. The aim of this procedure is two-fold: first, to give an overview on a topical theme within environmental economics; second, to give an introduction to some recent developments in non-convex economics.

2. Shallow lake dynamics

The ecological model that will be used in the following gives a very simplified representation of the complex ecological feedback mechanisms that are active in a shallow lake; essentially, it only captures the possible occurrence of a catastrophic shift in the state variable. The very stylizedness of the model is however also an advantage, as the conclusions that can be drawn from it will most likely be applicable to a wide range of other situations.

2.1 Ecological processes in a shallow lake

The following, very simplified, description of the ecological mechanisms acting in a shallow lake is based on the account in Scheffer's book⁷, to which the reader is directed for the full story.

A shallow lake, that is a lake whose depth is up to 3 m, is characterised by two facts: large water plants may grow all over the lake, and in summer it is, unlike deep lakes, not stratified in a warm top layer and a cold bottom layer. The second fact implies that interaction with the sediment may be important in a shallow lake; in particular, pollutant that is sedimented out may be resuspended in the water column. In the simplified shallow lake model, this effect is however entirely ignored.

The lake vegetation is determined by the nutrient level. When this is low, the vegetation consists mostly of small water plants. But artificial fertilisers used in agriculture are by rainfall washed off the fields and into the lake, raising the nutrient level. The primary effect is to make the growth of large water plants possible. Next, the phytoplankton biomass (algae) increases, making the lake turbid; also, the leaves of the water plants become covered by a periphyton layer. Both processes lead to a decrease of the amount of light reaching the plants, which eventually die.

As the water plants disappear, the food web in the lake changes radically. A natural predator of phytoplankton is zooplankton, which hides between the water plants during daytime. They lose their natural shelter and are increasingly predated on by fishes. This reinforces the growth of phytoplankton and increases the turbidity of the lake further. Moreover, the water plants are important in keeping the lake sediment intact; in their absence, the sediment is more easily resuspended into the water, again increasing the turbidity of the lake.

Now, even if the nutrient loading is reduced, the water plants will not be able to regrow easily, as the turbidity of the lake remains high. Nutrient reduction for long periods of time, or removal of fishes that predate on zooplankton may be necessary to restore the lake to a clean state again. In some cases the turbidity of the lake is however irreversible.

2.2 Mathematical model of a shallow lake

The simplified shallow lake dynamics are modelled as a single non-linear differential equation

$$s' = a - g(s) = a - \left(bs - \frac{s^2}{1 + s^2} \right).$$

Here $s = s(t)$ is the state variable of the system, that is, the concentration of phosphorus, one of the main nutrients, in the shallow lake. The variable $a = a(t)$ is the action that economic agents undertake: here it is the use of artificial fertilisers containing phosphorus on the fields surrounding the lake, which is then washed into the lake by rainfall, yielding a net inflow of

nutrients. The parameter b denotes the sedimentation rate at which phosphorus leaves the water column and enters the sediments at the bottom of the lake. In this very simple model, it is assumed that it remains in the sediment indefinitely: a more complex model would take into account that phosphorus from the sediments can feed back into the lake, especially if the vegetation of the lake has disappeared. Finally, the non-linear term models all the complex feedback mechanisms of the lake.

The curve $a = g(s)$ is called the *response curve* of the system. For large values of the sedimentation rate, typically for $b \geq 0.65$, the response is monotone, as illustrated in Figure 1.




Figure 1: Shallow lake dynamics, monotone case ($b=0.68$).

The curve indicates the locus of the steady states (a,s) of the shallow lake dynamics; all steady states are attracting.

The figure indicates the steady state s to which the state of the lake will tend for a given loading level a . The most important characteristic of this response is that small changes in a will imply small changes in s .

If the sedimentation rate is lower, the non-linearities of the system change this simple behaviour profoundly, as is illustrated in Figure 2.




Figure 2: Shallow lake dynamics, reversible case ($b=0.52$): effect of changing the loading slowly.

As in Figure 1, the thinly drawn curve, mostly overlaid by the thick curve, indicates the locus of the steady states (a,s) of the shallow lake dynamics; the dashed part of the curve corresponds to unstable equilibria. The thick curve shows the response of the state of the lake to slow changes of the loading. If the loading starts at $a=0$ and is increased, at first the amount of pollutant in the lake increases proportionally; but at $a=\bar{a}_{cr}$ it flips catastrophically to a much higher value of s , associated with a turbid lake. If then the loading is decreased again, the lake does not instantly jump back to a clean state; only if a is reduced to the lower value $a=\underline{a}_{cr}$, the lake flips back again to a low value of s , associated with a clean lake.

In this figure, two curves are shown. The thin curve, which is partly overlaid by the thick curve, gives all possible steady states of the system, that is, all combinations (a,s) of actions and states for which the rate of change of the state is zero. This curve consists of three branches. The upper and the lower branch correspond to stable steady states: those are states such that the system tends towards them if it is already sufficiently close. The middle branch consists of unstable steady states, from which the system drifts away. Steady states of the lower and upper branch are respectively referred to as low and high pollution steady states.

If we imagine the system to be initially in a state where there is no loading and no phosphorus in the lake, the thick curve shows the response of the system to a very slow increase of the phosphorus loading. Increasing the loading a from an initial value $a=0$, we move along the lower branch of the thick curve. Note that the phosphorus concentration s changes initially only very little as the loading increases. However, there is a critical point $a = \bar{a}_{cr}$, a tipping point, indicated by a black dot on the horizontal axis, where the low pollution steady state loses stability. The system will then move rapidly along the dashed trajectory to a high pollution steady state.

If we try to restore the system to its previous low pollution state by lowering the loading, we move along the upper branch of the thick curve for decreasing values of a .

The system will remain in the high pollution state until the loading is decreased quite significantly below the critical loading that made the system flip into the high pollution state. In the reversible case, which is illustrated by Figure 2, if the loading is decreased sufficiently much, the system encounters another tipping point $a = \underline{a}_{cr}$ at which it flips rapidly back to the low pollution steady state. As increasing or decreasing the external loading leads to different responses, the shallow lake system is said to exhibit hysteresis; this typically occurs for intermediate values of the sedimentation rate, that is, for $0.5 < b < 0.65$.

If the sedimentation rate is still lower, that is if $0 < b \leq 0.5$, then the system response is irreversible; this case is illustrated in Figure 3.




Figure 3: Shallow lake dynamics, irreversible case ($b=0.49$): effect of changing the loading slowly. Legend as in Figure 2. In the irreversible case, if the lake is in a high-pollution state, which corresponds in the figure to the upper branch of the thick curve, even by a reduction of the loading to the minimum value $a=0$, the lake cannot be brought back to a clean state.

In these plots, the location of the middle branch of unstable steady states also provides information: if for fixed loading a the state s of the system is such that the point (a, s) lies below the middle branch, then the system is attracted to the low pollution equilibrium, whereas if it is above that branch, it is attracted to the high pollution equilibrium. Note that the middle branch decreases as loading increases, meaning that as the loading increases, the basin of attraction of the low pollution steady state steadily declines, whereas the basin of the high pollution steady state grows. If the distance of the steady state to the boundary of the basin of attraction is small, the steady state is said to have little resilience, as only a small perturbation to the steady state suffices to move the system out of that basin, and move it then to the other steady state.

3. Quasi-static shallow lake economics

Environmental pollution is an, often inescapable, side effect of agricultural or industrial production activities. The public has an interest in a clean environment as well as in something to eat: there is here a basic economic trade-off. To determine the optimal policy, the two interests are quantified as term of a welfare function. The analysis below is considered with the utopian "social planner" scenario, where there is a benevolent and omniscient social planner that tries to maximise social welfare. Also the possibility is considered that there are several of these social planners sharing access to a shallow lake, and that each of them is concerned by maximising the social welfare of his community.

First we shall be concerned with the quasi-static context, where the loading level is to be fixed for all time. Later, the dynamic context is considered, where the loading level is allowed to vary as a function of time.

3.1 Quasi-static analysis

To model the choice of the loading level a , we assume that there is a social planner that tries to maximise some social welfare stream $W(a)$ that depends on the loading and the amount of pollutant in the lake, by choosing the constant loading level a optimally. For a given choice of a , the system is given time to adjust to a steady state $s = s(a)$, which is a solution of the equation $s'(t) = a - g(s) = 0$, or equivalently, of the equation

$$a = g(s).$$

Welfare is then evaluated at this steady state. Investigating the properties of these steady states constitutes a quasi-static analysis of the system, as neither dynamic adjustment effects are taken into account, nor the possibility that loading levels might change over time.

The social welfare stream W will be the sum of the utility U_A that affectors of the state of the lake, for instance farmers, obtain from their use of artificial fertiliser and the disutility D_E that enjoyers of the lake, tourists and fishermen, suffer as a result of pollution. In the economic shallow lake model, these utilities are modelled as $U_A = \log a$ and $D_E = cs^2$ respectively; the parameter c is a measure for the economic importance of the shallow lake, or the economic costs of pollution. This yields the following expression for social welfare

$$W = U_A - D_E = \log a - cs^2$$

Note that these are really utilities from different groups of lake users; the planner tries to determine the socially optimal use of the lake. As mentioned, in this section a quasi-static approach is taken: the social planner is required to find the loading level a such that in a corresponding steady state of the shallow lake system, social welfare is optimal.

In order to do this, steady state values $s = s_L(a)$ for low pollution steady states and $s = s_H(a)$ for high pollution steady states are solved from the steady state equation $s'(t) = 0$. These are the functions whose graphs trace out the two stable steady state loci in Figure 2. Using these solutions, welfare in low and high pollution steady states, respectively W_L and W_H , can be derived as functions of the loading a , yielding for $i = L$ or $i = H$ that

$$W_i(a) = U_A(a) - D_{E,i}(a) = \log a - cs_i(a)^2.$$

Note that $D_{E,L}(a)$ and $D_{E,H}(a)$ indicate respectively the disutility from pollution in the low pollution and high pollution steady state associated to the loading level a . The social welfare is then the maximum of the low pollution and high pollution welfare functions

$$W(a) = \max\{W_L(a), W_H(a)\}.$$

At a maximum of W , the marginal welfare from increasing the loading should vanish, that is $W'(a) = 0$, or, equivalently, the marginal utility of the affectors should equal the marginal disutility for the enjoyers. This is the condition for economic equilibrium:

$$U'_A(a) = D'_{E,i}(a).$$

Performing the differentiations yields

$$\frac{1}{a} = 2c s_i(a) s'_i(a).$$

The value of $s'_i(a)$ is obtained by differentiating the steady state equation $a = g(s_i(a))$ with respect to a . Solving for $s'_i(a)$ yields

$$s'_i(a) = \frac{1}{g'(s_i(a))}.$$

Substitution into the economic equilibrium equation yields finally

$$\frac{1}{a} = \frac{2c s_i(a)}{g'(s_i(a))}.$$

Alternatively, this condition can be obtained by using a Lagrange function approach, where the social planner maximises welfare W subject to the condition that $a - g(s) = 0$. One major advantage of this approach is that no distinction has to be made between separate solution branches of the equilibrium equation $a - g(s) = 0$. The Lagrange function for this problem reads as

$$L = \log a - c s^2 + p (a - g(s)).$$

Note the appearance of a new variable p , which will turn out to be the shadow value of the lake; that is the price the social planner is willing to pay for adding another unit of pollutant. As the planner is never willing to pay for more pollution, the shadow value is going to be negative.

If the loading-state pair (a, s) maximises social welfare W subject to the condition $a - g(s) = 0$ that the lake is in dynamic equilibrium, then the triple (a, s, p) satisfies necessarily the Lagrange conditions

$$0 = \frac{\partial L}{\partial a} = \frac{1}{a} + p,$$

$$0 = \frac{\partial L}{\partial p} = a - g(s),$$

$$0 = -\frac{\partial L}{\partial s} = 2c s + p g'(s).$$

Solving p and a from the first two equations as

$$a = g(s), \quad p = -\frac{1}{a} = -\frac{1}{g(s)},$$

and substitution of these values into the third equation yields the relation

$$\frac{1}{a} = \frac{2c s}{g'(s)}$$

obtained previously. Besides reconfirming the result, we have also found an interpretation of the Lagrange conditions: the first, $\partial L/\partial a = 0$, is the economic equilibrium condition for the loading: the marginal increase in benefits from polluting the lake should equal the marginal shadow cost of that pollution. The second, $\partial L/\partial p = 0$, is the dynamic equilibrium of the lake, and the third condition $\partial L/\partial s = 0$ is the economic equilibrium equation for the pollutant in the lake: the marginal welfare cost of increasing the amount of pollutant should equal its marginal benefits.




Figure 4: Marginal utilities ($b=0.52$, $c=1.3$).

At an economic equilibrium, the marginal utility $U'_A(a)$ to the affectors from adding another unit of loading, indicated by the dashed graph, equals the marginal disutility $D'_{E,i}$ of the enjoyers, solid graphs, where $i=L$ or $i=H$ according to whether the lake is in the low or in the high pollution regime. Note that the disutility in the low pollution regime increases sharply as a approaches the critical value, forcing at least one solution of the economic equilibrium equation to be on the low pollution branch. The two intersections marked with a dot correspond to local maxima of the welfare function.

We show these marginal utilities in Figure 4 as functions of the loading level a . The marginal utility for using artificial fertiliser, indicated by the dashed graph, decreases with an increase of the use of fertilisers: this is the usual law of diminishing returns. The marginal disutilities in the low and high-pollution states respectively are indicated by solid graphs in the figure. The effect of increasing the input level of fertiliser on the disutility generated in the low pollution steady state also follows a standard pattern: the marginal disutility in the low pollution equilibrium increases with the input level. Note that it has a vertical asymptote at the upward tipping point: this is a consequence of the fact that already a little below the tipping point, small changes in the input level induce large changes in the pollution level of the lake. It is this feature of tipping points on which the rising variance early warning systems for ecological transitions are based^{8,9}.

The effect of the pollution input level on the high pollution steady state is at first sight a little surprising, as the marginal disutility from pollution in that state decreases instead of increases with increasing input levels. Partly this can be understood as before: slightly above the downward tipping point, small increases of loading imply large increases in the phosphorus concentration of the lake, and this response diminishes as the loading increases further.

Turning to the determination of the (local) maxima of the welfare function, note that these occur at those intersection points of the graphs of marginal utility and marginal

disutility where marginal utility from increasing the loading is initially larger than marginal disutility; the other intersection points correspond to local minima. In Figure 4 we have two such intersection points, at $a \approx 0.047$ (high pollution branch) and $a \approx 0.072$ (low pollution branch) respectively. In this case, it is immediately obvious that the latter intersection corresponds to a global maximum of the welfare function, as pollution is lower and at the same time the use of fertiliser is higher compared to the first intersection point. This is borne out by the graph of the welfare function $W(a)$, which is given in Figure 5.





Figure 5: Welfare ($b=0.52$, $c=1.3$).

The graphs of the full welfare functions $W_L(a)$ and $W_H(a)$ (dashed), as well as the function $W(a)=\max\{W_L(a),W_H(a)\}$ (solid). Welfare drops significantly as the loading a is increased past the tipping point value. Note also how the maximal value of welfare is obtained for a loading level that is very close to the critical level.

In this figure, total welfare is drawn as a solid graph, and the welfare in the high pollution branch is indicated by a dashed graph. Notice that at the tipping point, welfare drops instantly by a large amount: of course, this is what we mean by an environmental catastrophe. What is interesting though is that the welfare maximising loading level is very close to the tipping point. This is fairly typical for maximising behaviour, to strain a system towards the "maximal sustainable yield". The danger of this approach is that if the model of the decision maker is not quite correct, or if the state of the system is subjected to small random shocks that are neglected in the modelling stage, the real system may easily cross the tipping point and flip to the high-pollution, low-welfare state.

Other situations are possible as well; one of them is illustrated in Figure 6.



Marginal utilities and welfare (b=0.63, c=0.35).

Legend as in Figures 4 and 5. Again, there is a huge drop in welfare at the tipping point; however, as the economic importance c of the ecosystem to the social planner is very low, maximal welfare is obtained for a high loading, high pollution level.

Here again we see three points of equilibrium of marginal utility and disutility, two of which correspond to local maxima of the welfare function. Examination of the total welfare function shows, for increasing a , first the familiar pattern of welfare taking a local maximum very close to the tipping point. But then the welfare function rises again, until it reaches, for a much higher loading level, its global maximum. In this case, it turns out, the economic service the lake provides as a waste dump of phosphorus outweighs the economic importance of the ecological services.




Figure 7: Marginal utilities and welfare (b=0.49, c=0.8).

Legend as in Figures 4 and 5. The value of b is below 0.5, implying that the lake is irreversible: if the lake is initially clean, but the loading is allowed to grow past the critical value, the ensuing welfare loss is irrecoverable.

The situation is different again if the lake is irreversible; this is illustrated in Figure 7. Note that the welfare function looks much like that of the reversible case of Figure 5. There is one important difference though: if the loading a is allowed to increase above the critical value, the lake flips, and the welfare level drops catastrophically. What is more, it cannot be restored to its former value: the welfare loss is irrevocable.

We find that we can classify the different situations, described by different values of the parameters b and c , by determining whether the high pollution or the low pollution furnishes a global maximum of the social welfare function. The result of this classification is shown in Figure 8.




Figure 8: Effect of welfare optimisation on the lake.

If the pair (b, c) corresponds to a point to the right of the dashed curve, the welfare function is maximal at a unique loading level, which moreover can always be reached. To the left of the dashed curve, there are two candidate maximisers. In the 'high pollution' region, maximal welfare is obtained in a high loading, high pollution state of the lake; in the 'low pollution' region, welfare is maximised in a low loading, low pollution state. For pairs (b, c) in the third region, the decision maker has to take care not to let the loading increase past the critical level, on penalty of suffering a permanent irrecoverable welfare loss.

There are several regions in the plot shown in Figure 8. The region to the right of the dashed curve is the region where the social welfare function has a single local maximum, which is at the same time a global maximum. To the right of the dashed curve, the social welfare function has two separate local maxima. The region where this is the case is divided into three subregions: the high pollution region, the low pollution region, and the region labelled "dependent on initial state". The significance of the first two is clear: if the pair (b, c) takes a value in the low pollution region, the low pollution equilibrium furnishes the global maximum of the social welfare function; the analogous result holds for the high pollution region. Moreover, the social planner can always reach the global maximum by adjusting the loading a appropriately.

Matters are different in the third region: for parameter values in this region, the low pollution equilibrium furnishes the global maximum of the social welfare function, but this maximum cannot be reached if the lake is allowed to pass beyond the tipping point, as it cannot then be restored to a low pollution state. This is the region where wrong decisions, or mere inattention of decision makers, can lead to an ecological catastrophe that has far-reaching welfare effects.

3.2 The quasi-static shallow lake game

The fate of shallow lakes, and natural resources in general, are often the outcome of a game that several policy makers are playing: a lake might be bordered on by several communes or several countries. The lake services are often best modelled as a public good, and as such they are subjected to the tragedy of the commons. The pollutant in the game is a stock variable, generating costs for all lake users, whereas the individual loadings are flow variables, generating individual benefits. If the lake users do not cooperate, they have a

tendency to emphasise their personal benefits over the commonly shared costs: the short-time flow effects tend to dominate the long-time stock effects.

The quasi-static shallow lake game, which will be discussed now, is briefly mentioned by Mäler *et al*¹. Let us examine the situation more closely. Assume that there are n agents; the welfare function of agent i is given as

$$W^i = U^i_A - D^i_E = \log a_i - c_i s^2.$$

Here a_i is the loading level of agent i , and c_i is his cost parameter. The total loading is now the sum of the individual loadings; we therefore obtain the following equation for the steady state of the lake:

$$0 = a_1 + a_2 + \dots + a_n - g(s),$$

where, as always, $g(s) = bs - s^2 / (s^2 + 1)$. The agents play a game: that is, each agent tries to maximise his individual payoff W^i by choosing his loading level a_i appropriately.

As an aside, it should be mentioned that this setup is almost equivalent to a standard public goods game; the author is indebted to P. Heijnen for this observation. The equivalence is effected by solving $s = s(a_1 + \dots + a_n)$, where $s(a) = s_L(a)$ or $s(a) = s_H(a)$, from the steady state equation. Introducing the equivalent payoff function $\tilde{W}^i = e^{W^i}$ leads then to considering

$$\tilde{W}^i = e^{-c_i s (a_1 + \dots + a_n)^2} a_i = f(a_1 + \dots + a_n) a_i.$$

The main difference with the standard case is that the function s , and consequently the function f , can be multi-valued.

We consider non-cooperative play, and we look for Nash equilibria of this game: that is, we look for an n -tuple (a^*_1, \dots, a^*_n) of actions, which is such that, given the action $a^*_1, \dots, a^*_{i-1}, a^*_{i+1}, \dots, a^*_n$ of the other agents, agent i cannot do better than to choose action a^*_i himself. Assuming the actions of the other players known - we drop the stars - agent i determines his own action by solving the constrained optimisation problem to maximise his welfare, given the steady state equation.

As above, it is convenient to formulate the maximisation problem of the agent as a problem with an equality restriction; that is, the dependence of the state on the action $s = s(a)$ is not solved from the equilibrium equation $a = g(s)$, but this equation is kept in the formulation of the maximisation problem. The Lagrange function reads as

$$L_i = \log a_i - c_i s^2 + p_i \left(\sum a_j - g(s) \right);$$

here p_i is the shadow cost which agent i assigns to the state of the lake, that is, the marginal cost associated to a unit increase of the amount of pollutant in the lake. The action of the agent is a solution of the necessary conditions of this problem

$$\frac{1}{a_i} = p_i, \quad p_i g'(s) = 2c_i s, \quad \sum_{j=1}^n a_j = g(s).$$

If we eliminate p_i from the first two equations, we obtain

$$\frac{1}{a_i} = \frac{2c_i s}{g'(s)}.$$

We recognise the economic equilibrium equation; but here we have this equation for every $i = 1, \dots, n$. To obtain the Nash equilibria steady states, these n equations have to be combined with the steady state equation

$$\sum_{j=1}^n a_j - g(s) = 0.$$

Eliminating the variables a_i from these equations, we obtain the following expression for the equilibrium states

$$\frac{1}{g(s)} = \frac{2c_{NC}s}{g'(s)},$$

where

$$\frac{1}{c_{NC}} = \frac{1}{c_1} + \dots + \frac{1}{c_n}.$$

The point of these equations is that the coefficient c_{NC} is the only quantity that depends on the number and the type of the players. There are two important special cases: if there is only one player, then $c_{NC} = c$. In the symmetric case of n identical players, that is, if $c_1 = \dots = c_n = c$, we find that $c_{NC} = c/n$.

To illustrate what effects of adding a player has on the Nash equilibrium pollution levels of the lake, the symmetric case is considered, for which $a_1 = \dots = a_n = a$, and for $n = 1$ and $n = 2$ the marginal utility

$$\frac{dU_A^i}{da} = \frac{n}{g(s)}$$

and the marginal disutilities

$$\frac{dD_{E,k}^i}{da} = \frac{2cs}{g'(s)},$$

where $k=L,H$, are given in Figure 9 as functions, not of the actions a , but of the state s .

marginal welfare




Figure 9: The effect of adding a player ($b=0.52$, $c=1$).

Marginal utility (dashed curve) and disutility (drawn curve) of increasing the loading, expressed as a function of the state s , for the case of one and two countries bordering the lake. There are in both cases three intersections of the two curves; only the leftmost (clean) and the rightmost (polluted) intersection are potential Nash equilibria. The location of the clean intersection hardly changes as the number of player increases, since the marginal disutility rises sharply shortly before the flipping point. In contrast to this, the location of the polluted intersection shifts significantly values of higher pollution as the number of players increases.

The marginal disutility for low pollution states vanishes at $s = 0$, and it has an asymptote at the first tipping point, where $g'(s) = 0$. Therefore, its graph always crosses the graph of the marginal utility in the low pollution region. We can conclude that there is always a low pollution candidate location for a Nash equilibrium for the quasi-static game. Moreover, we can readily conjecture from the figure that the low pollution Nash equilibria will be located closer and closer to the tipping point as the number of players increases.

In the figure, we see that there are two crossings of the graphs of marginal utility and marginal disutility in the high pollution region. As before, the lower of the two corresponds to a local minimum of the welfare function of the players, and hence has no economic significances. The other crossing determines the location of a high pollution candidate Nash equilibrium. We see that the corresponding state of the system increases as the number of players increases.

However, to conclude that the crossings determine the location of true Nash equilibria values of the state would be premature. It still has to be checked whether, given that the actions of the other players correspond to the value determined by the crossing, the candidate value for the action does indeed maximise utility for a given agent. The difficulty is exactly the same as the distinction between local and global maxima in the simpler social planner problem discussed above. It turns out that the difficulty is not always hypothetical: there are parameter sets where either the low pollution or the high pollution candidate equilibrium does not furnish a maximum of the individual payoff. In such a situation, a player obtains a higher payoff by deviating from the symmetric candidate strategy, and the crossing does not furnish a Nash equilibrium. If the lake is irreversible, there is the additional complication that some states of the lake cannot be reached, depending on the initial state of the lake.




Figure 10: Classification of Nash equilibria configurations ($n=2$).

For the two-player quasi-static shallow lake game, this diagram summarises the possible outcomes. To the right of the dashed curve, there is only a single candidate Nash equilibrium, which also furnishes a true Nash equilibrium. To the left of the curve, there are always two candidates. In the region 'clean NE' and 'polluted NE', respectively the polluted and the clean candidate equilibria fail to be Nash. In the region '#NE=2', both candidates are Nash; however, the value of the payoff is always higher in the clean Nash equilibrium. In the region '#NE=1 or 2', also both candidates are Nash equilibria; however, if the initial pollution is too high, then the clean Nash equilibrium cannot be reached. Also in this region, the payoff in the clean Nash equilibrium is higher than in the polluted one. Finally, in the part of the region 'polluted NE' above the dotted line, the payoff in the clean candidate Nash equilibrium, which fails to be truly Nash, is higher than in the polluted Nash equilibrium: in this region, the prisoners dilemma mechanism is at work.

As noted, the principal property of the solution is that the high pollution loading increases significantly if there are more and more players. This deterioration of the outcome reflects a fundamental distinction between the pollutant s in the lake, which is a common stock variable, and the phosphorus loadings a_i , which are individual flow variables. If the number of players increases, the relative effect of individual actions on the state variable decreases. Moreover, conservative behaviour as restricting the phosphorus loading is not rewarded, as other agents can freeride on this by increasing their loading. This is an example of the tragedy of the commons.

However, in the low pollution region, the tragedy is initially mitigated by the fact that at the tipping point, the marginal disutility of adding pollutant is infinite: more and more players do push the low pollution candidate location closer and closer to the brink, but not over it. Of course, if there are sufficiently many players, the low pollution location will not furnish a welfare maximum, given the actions of the other players and players will deviate to the high pollution outcome.

The situation should be contrasted with the welfare gains that can be obtained if all agents cooperate. In that case, we can consider a single joint welfare function

$$W = \sum_{i=1}^n (\log a_i - c_i s^2).$$

Again by solving the constrained optimisation problem, we find that the state has to satisfy a state equilibrium condition of the form

$$g'(s) - 2c_C s g(s) = 0,$$

where now

$$c_C = \frac{1}{n} \sum_{i=1}^n c_i.$$

Note that in the symmetric case $c_C=c$, and that the outcome is independent of the number of players: the eventual steady state reached is the same each player would reach if he would have been a single social planner.

3.3 The repeated game

Brock and de Zeeuw extended the quasi-static analysis of the shallow lake game by considering an infinitely repeated game¹⁰. They focussed on the question whether the cooperative solution of the quasi-static game can be supported by a trigger strategy of the players.

The idea is the following: at the beginning of the repeated game, all players state their intention to follow a cooperative strategy, which will lead to a high level of welfare. However, this situation is not stable in the sense that if all players but one follow the cooperative strategy, whereas if one player deviates from it by choosing his action non-cooperatively the payoff of the deviating player will typically be higher than that which he would have obtained if had followed the cooperative action. By symmetry, all players have therefore an incentive to deviate from the cooperative strategy, and they will in the end all play non-cooperative strategies. The social welfare in this case will be lower than in the cooperative case.

To stabilise the cooperative solution, players can announce to play a trigger strategy, announcing that they will play the cooperative strategy unless one player deviates. After that event, the player playing the trigger strategy will play the Nash equilibrium strategy with the lowest payoff. An agent will have an incentive to deviate if he does not care greatly for future welfare benefits. To evaluate the trigger strategy, we have to know the time preferences of the agents; these are expressed by a discount factor δ at which a contribution W to the social welfare tomorrow is equally welcome as a contribution δW today. It is obvious that if the discount factor is equal to 0, that is, if agents do not care about the future at all, a trigger strategy will have no bite, and every agent will deviate. At the other end of the spectrum, if the discount factor is 1, welfare benefits in the distant future weigh as heavily as present benefits, and deviations will never be optimal. The effectiveness of the trigger strategy can therefore be expressed quantitatively by the minimal discount factor δ_{min} for which no player has an incentive to deviate. Let W_d denote the welfare gain in the one-shot game for a player that deviates, while the other players cooperate, let W_{Nash} denote the welfare gain in the Nash equilibrium and let W_c denote the welfare gain under cooperation. Brock and de Zeeuw give the relation

$$\delta_{min} = \frac{W_d - W_c}{W_d - W_{Nash}}.$$

The dependence of the minimal discount factor δ_{min} on the parameter b and the number of players n is illustrated in Figure 11.




Figure 11: Contours of minimal discount factor for 'worst Nash' trigger ($c=2$).

The plot illustrates the minimal discount factor δ_{min} for which a trigger strategy, threatening to switch to the most polluted Nash equilibrium, still stabilises the cooperative outcome. Its dependence is shown on the sedimentation rate b , as well as the number of players n , treated as if the number of players were a continuous variable; only integer values of n have any real meaning. The minimal discount factor is discontinuous along the ridge which runs from $(b,n)\approx(0.5,2.7)$ to $(b,n)\approx(0.8,4)$. The ridge is associated with the appearance of a low payoff polluted Nash equilibrium. Note that the threat works well for large values of n , that is, δ_{min} is very low there, since there the Nash outcome is very bad.

In the figure, contour lines of constant values of δ_{min} are shown. They are computed for a continuum of values of n ; obviously, only the integer values of n are significant. Note the thick and rather wiggly curve, running from $(b,n)=(0.5,2.7)$ upwards. At this curve, the value of δ_{min} is discontinuous: for instance, for $b=0.6$, it changes from approximately 0.67 before the discontinuity to 0.03 after it. This reflects the fact, already noted in the previous subsection, that increasing the number of players n has the roughly same effect as decreasing the economic weight c of the lake. In this case, at the discontinuity ridge, a high pollution Nash equilibrium springs into existence, and the strength of the threat, to switch to the Nash equilibrium with the lowest pay-off, is suddenly much enhanced.

Of course, it is debatable whether a player that wishes to punish a deviation from cooperative behaviour will actually execute a Nash equilibrium strategy with minimal pay-off, as the cost to himself are substantial. Therefore, we should contrast the situation of Figure 11 to the situation that after a deviation again Nash strategies will be played, but now those with the highest pay-off. The corresponding values of δ_{min} are given in Figure 12.




Figure 12: Contours of minimal discount factor for 'highest payoff Nash' trigger ($c=2$).

Legend as in Figure 11, only that the threat is to change to the Nash strategy with the highest payoff. Note that the discontinuity ridge of Figure 11 disappears, but that there appears a different discontinuity ridge, now associated with the disappearance of a high payoff clean Nash equilibrium.

The discontinuity ridge of Figure 11 has disappeared; a different ridge appears, but for much larger values of b and n . We conclude that as the threat loses its "bite", the trigger strategy is less effective in stabilising cooperative behaviour. Still, the values of δ_{min} are still moderate: a value of $\delta_{min}=0.7$ corresponds to an effective time horizon of about 7 periods.

4. Dynamic shallow lake economics

In the quasi-static context, which we considered in the previous section, there is for a large range of parameters a low pollution steady state, or in the game context, a low pollution Nash equilibrium; this is largely due to the fact that in this context, agents and players are committed to strategies which cannot change in time. Therefore, the sharp increase in marginal disutility at the tipping point keeps the solutions in the low pollution region.

These results, rather positive from an environmental point of view, are weakened if a full dynamic context is considered, as we shall do in the present section. First we consider a single social planner that allows the amount of phosphorus loaded into the lake to change as a function of time: $a=a(t)$. The amount $s=s(t)$ of phosphorus in the lake still evolves according to the evolution law

$$s' = a - bs + \frac{s^2}{s^2 + 1}.$$

Social welfare W is now taken to depend on the total welfare stream $w(t) = \log(a(t)) - c s(t)^2$ for t ranging from 0 to infinity; we consider

$$W = \int_0^{\infty} e^{-\rho t} w(t) dt = \int_0^{\infty} e^{-\rho t} (\log a(t) - c s(t)^2) dt.$$

That is, we consider the case that future contributions to the welfare stream are discounted at a discount rate ρ . To solve this optimisation problem, we form the current value Pontryagin function (also called pre-Hamilton or unmaximised Hamilton function)

$$P = \log a - c s^2 + p (a - g(s)).$$

Note the similarity to the Lagrange function: in fact, the only difference is that the state $s=s(t)$, the loading $a=a(t)$, and the shadow value $p(t)$ of the lake are all three functions of time. Pontryagin's maximum principle now states that a solution $(a(t), s(t), p(t))$ of the problem to maximise W , taking into account the evolution law of the pollutant in the lake, necessarily satisfies the following three conditions:

$$0 = \frac{\partial P}{\partial a} = \frac{1}{a} + p,$$

$$s' = \frac{\partial P}{\partial p} = a - g(s),$$

$$p' = \rho p - \frac{\partial P}{\partial s} = 2cs + (\rho + g'(s))p.$$

Again, the formal similarities to the Lagrange conditions should be noted. After elimination of a , we retain a system of two differential equations. In order to obtain a solution to the system, we need to specify two additional conditions. A natural condition is putting $s(0) = s_0$, where s_0 is the initial state of the lake. The second condition, the so-called transversality condition, is more involved; it states that the discounted shadow value of the lake at the end of the time period vanishes:

$$\lim_{t \rightarrow \infty} e^{-\rho t} p(t) = 0.$$

In practice, solving a problem of this type boils down to first finding steady states (\bar{s}, \bar{p}) of the system of differential equations, for if $p(t) = \bar{p}$ is constant for all t , then the transversality condition is surely satisfied. The second step is then to determine solution trajectories of the system that tend in the limit $t \rightarrow \infty$ to the given steady state. This procedure is iterated for all steady states.

In the shallow lake dynamic optimisation problem, there are either one or two steady states that can be limits to solution trajectories. If two such states exist, there may be initial states s_0 for which there are several candidate solution trajectories starting at s_0 . The trajectory with the highest value of the welfare then furnishes a solution to the maximisation problem^{1,11}. A particular feature of the shallow lake problem is that there may be a single initial state s_* at which two candidate solution trajectories start that generate the same welfare, and which both yield a true solution to the problem. Such an initial state is an *indifference state*, as the planner is indifferent between two qualitatively different loading schedules. Note that other names for indifference points are "Skiba points", "Dechert-Nishimura-Skiba points" or "Dechert-Nishimura-Skiba-Sethi points" in the economic literature, and "shock points" in the optimal control literature.

Solution trajectories are traditionally shown in an (s, a) -diagram instead of an (s, p) -diagram; the transformation is performed by making use of the relation $a = -1/p$. A typical solution of the shallow lake problem is given in Figure 13.




Figure 13: Typical solution of the dynamic shallow lake problem ($b=0.55$, $c=0.5$, $\rho=0.03$).

Two solution trajectories are shown (thin curves) that originate from the focus in the middle of the plot and that tend either to the left or the right steady state. The thick curves indicate the parts of the trajectories maximising welfare. At $s \approx 0.85$ there is an indifference state, at which the decision maker is indifferent between the low loading, low pollution solution leading to the left, and the high loading, high pollution solution leading to the right.

The values of b , c and ρ are respectively 0.55, 0.5 and 0.03. The steady states are indicated by dots; they are located at approximately $(\bar{s}, \bar{a}) = (0.35, 0.08)$ and $(\bar{s}, \bar{a}) = (1.76, 0.21)$. The curves through these steady states are the solution curves of the system of differential equations that tend to these states in the limit $t \rightarrow \infty$.

Note that if $0.58 \leq s_0 \leq 1.21$, there are several possible choices for the initial loading level a_0 and the associated shadow value p_0 . The initial values that maximise welfare are those points that are located on the thickly drawn curves. There is a single point $s_* = 0.83$ for which two different pairs of initial values (s_0, a_0) lead to solution trajectories that optimise social welfare: one is associated with initially very little, though gradually increasing, loading, tending to the low pollution steady state; the other is characterised by overall high loading, and that solution tends to a high pollution steady state.

The result of the optimisation analysis can be given as a loading policy $a = a(s)$, combined with the implied steady states and their basin of attractions - that is, the set of initial states which will be steered by the optimal loading policy to a given steady state. For the situation of Figure 13, these data are given in Figure 14.




Figure 14: Policy function and optimal dynamics ($b=0.55$, $c=0.5$, $\rho=0.03$)

Retaining only the relevant parts of the solution trajectories shown in Figure 13 gives the optimal policy function $a=a(s)$; note the discontinuity at the indifference point. On the horizontal axis, the resulting dynamics are given. The indifference point is marked by the square; the two points indicate steady states of the system under optimal management.

In the next subsections, we shall discuss how these solutions depend on the problem parameters.

4.1 In the footsteps of Ramsey

Ramsey famously asserted in his 1928 article about optimal savings that to take any other discount rate than $\rho = 0$ is "ethically indefensible". In the present context, he probably would have maintained that the functional

$$\int_0^{\infty} (\log a(t) - c s(t)^2) dt$$

should be maximised. There is an interpretation problem here, though: in most cases, the value of the integral is either plus or minus infinity, and there is then no good way to compare the values. There are sophisticated mathematical concepts for dealing with this problem (overtaking optimality, catching-up optimality), but for the purpose of this paper, it is sufficient to think of the average welfare flow instead, which is given as

$$\bar{w} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T (\log a(t) - c s(t)^2) dt.$$

This quantity is usually finite; moreover, it is clear that if loading $a(t)$ and pollutant $s(t)$ tends to constant levels \bar{a} and \bar{s} respectively, the average welfare flow will be equal to

$$\bar{w} = \log \bar{a} - c \bar{s}^2.$$

Looking at the three conditions obtained from the Pontryagin maximum principle, we see that the steady states $(s(t), a(t), p(t)) = (\bar{s}, \bar{a}, \bar{p})$ of that system satisfy the conditions

$$0 = \frac{1}{\bar{a}} + \bar{p}, \quad 0 = \bar{a} - g(\bar{s}), \quad 0 = 2c\bar{s} + \bar{p}g'(\bar{s}).$$

These conditions are identical to those we have obtained in the quasi-static case. As the welfare flow here is also equal to the welfare function in that case, we conclude that the Ramsey and the quasi-static context are equivalent. In particular, Figure 8 summarises the qualitative properties of the optimal solution depending on the parameters; but in the Ramsey context, this figure gives information about the solution structure of a dynamic problem, which moreover can be expected to be the limiting case of the problem with positive discounting.

4.2 Positive discounting

Policy practice has not followed Ramsey's categorical imperative: in almost all policy decisions, the wishes of the present generation are allowed to weigh more heavily than the projected interests of future generations.

The value of ecosystem services are to a large extent a given of a particular situation. For instance, Hein *et al.* analysed the values of the ecosystem services of De Wieden wetlands in the Netherlands¹², and Hein provided an estimate of the cost of removing pollutant from those wetlands¹³.

A policy maker that has to determine an optimal pollutant loading schedule has to accord some kind of weight to the future damages arising from pollution, and this is often a matter of choice. In the Netherlands, there used to be a standard future discount rate of 7% per annum; recently, this has been changed to 2.5% + $x\%$, where x is to be determined according to the special requirements of the valuation at hand.

Wagener¹¹ and Kiseleva and Wagener¹⁴ obtained a classification, analogous to that of Figure 8 of the quasi-static and the Ramsey case, for the fully dynamic case where the discount rate $\rho > 0$. The analysis is based on bifurcation theory, as optimal loading policies can be classified according to qualitative characteristics. In the discounted shallow lake problem, there are only three types, given in Figure 15.




Figure 15: Three classes of optimal policy functions

Top left: two attracting steady states, separated by an indifference point. Top right: two attracting steady states, separated by a repeller. Bottom: a single steady state. The distinction on the bottom line between a clean steady state (left) and a polluted one (right) is not qualitative but quantitative, and gives therefore *per se* not rise to separate classes.

A loading policy gives rise to either one or two long term attracting steady states; in the first eventuality, and if the sedimentation rate b is not too large, the steady state can be typified as being either low pollution or high pollution: however, the latter classification is not qualitative but quantitative.

If there are two attracting steady states, the policy function is either continuous, or it has a jump. If it is continuous, in between the two attracting steady states, there is also a repelling steady state. For initial states close to such a repeller, the optimal loading policy is to keep the loading at the repeller's level for a long time, before steering the lake either to a low or a high pollution equilibrium. Below we shall see that this situation typically occurs if the discount rate is large. The result can be formulated as follows: if the social planner is mainly concerned with the short run, there is a critical region of initial states of the lake at which the optimal policy of the planner is to defer the decision between an industry-friendly or an environmentally-friendly policy to a later date.

The situation is quite different if the optimal policy function has a jump; this commonly occurs if the discount rate is low, and if the planner is consequently mainly concerned with the short run. In that constellation, the planner wants the lake to leave the critical region around the jump point as quickly as possible, and he makes his choice between the two types of policy right away.

The dependence of the different possibilities on the parameters b and c are given in Figure 16; this figure should be compared to Figure 8.




Figure 16: Effect of dynamic welfare optimisation on the lake (top) and detail (bottom) for $\rho = 0.03$.

If the pair (b, c) corresponds to a point to the right and below of the drawn curves, there is only a single attracting steady state and the optimal policy function is continuous. For high values of c , the steady state corresponds to a clean lake, for low values of c , it corresponds to a polluted lake. In the region on the left, which is enclosed by the two drawn curves, there are two attracting steady states, separated by an indifference point. Moreover, in the detailed figure, close to the point C, there is a region where there are two attracting steady states, separated by a repeller, and where the optimal policy function is continuous.

5. Conclusion

In this review, outcomes of the shallow lake long-term interest conflict in a number of different settings have been presented, in particular in the context of quasi-static and dynamic social planning, of quasi-static non-cooperative play, and of the effect of trigger strategies in repeated quasi-static play. A characteristic feature of this interest conflicts, and of pollution problems in general, is the qualitative dichotomy in possible outcomes: the lake (or the ecosystem, or the climate) ends up in either a clean or in a polluted state, both of which, if attained, is stabilised by some kind of feedback mechanism. The social choice therefore always incorporates a qualitative aspect: the decision maker has to decide for or against production, for or against conserving the ecosystem.

This qualitative distinction between the possible socially optimal outcomes enables us to present the outcome of the analyses in the form of a bifurcation diagram (Figures 8, 10 and 16), which gives a graphical overview of the qualitative characteristics of the solutions, depending on the parameters of the problem. In the case of the repeated game, the outcome was a parameter diagram featuring a catastrophic set, where the economic quantity of interest changes discontinuously. The most critical region in these bifurcation diagrams is the "history dependent" region: in these cases, neglect by an actual decision maker that allows the ecosystem to flip can lead to large irrecoverable welfare losses.

It is readily apparent that this line of research, even in the case of the shallow lake problem, is far from being completed: for instance, the dynamic open loop¹ and closed loop^{15,16} games have still to be analysed in a similar fashion. It can be readily conjectured that the classification of dynamic Nash equilibria for small values of the discount factor ρ should resemble Figure 10, much as Figure 16 resembles Figure 8. Also the situation of asymmetric players is still to be investigated, as well as that of a shallow lake subject to stochastic dynamics¹⁷. Another interesting development away from the social planner context is the analysis of the political economy of the decision making¹⁸.

Acknowledgements

I would like to thank Pim Heijnen, Tatiana Kiseleva, Saeed Mohammadian Moghayer, Colin Rowat and Aart de Zeeuw for helpful comments and suggestions. The plots given in Figure 16 have been prepared by Tatiana Kiseleva. The support of a MaGW-VIDI grant of the Netherlands Organisation for Science (NWO) is gratefully acknowledged.

References

1. Mäler, K.-G., A. Xepapadeas & A. de Zeeuw. 2003. The economics of shallow lakes. *Environmental and Resource Economics* 26: 603–624.
2. Carpenter, S. R., D. Ludwig, and W. A. Brock. 1999. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications* 9:751–771.
3. Brock, W.A. & D. Starrett. 2003. Managing systems with non-convex positive feedback. *Environmental and Resource Economics* 26:575–602.
4. Tahvonen, O. & S. Salo. 1996. Nonconvexities in optimal pollution accumulation. *Journal of Environmental Economics and Management* 31:160-177.
5. Holling, C. S. 1973. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics* 4:1–23.
6. Scheffer, M. et al. 2001. Catastrophic shifts in ecosystems. *Nature* 413:591-596
7. Scheffer, M. 1998. *Ecology of shallow lakes*. Chapman & Hall. London.
8. Brock, W. A., and S. R. Carpenter. 2006. Variance as a leading indicator of regime shift in ecosystem services. *Ecology and Society* 11:2.
9. Carpenter, S. R., and W. A. Brock. 2006. Rising variance: a leading indicator of ecological transition. *Ecology Letters* 9:308–315.
10. Brock, W.A. & A. de Zeeuw. 2002. The repeated lake game. *Economics Letters* 76: 109-114.
11. Wagener, F.O.O. 2003. Skiba points and heteroclinic bifurcations, with applications to the shallow lake system. *Journal of Economic Dynamics & Control* 27:1533–1561
12. Hein, L. et al. 2006. Spatial scales, stakeholders and the valuation of ecosystem services. *Ecological Economics* 57:209–228
13. Hein, L. 2006. Cost-efficient eutrophication control in a shallow lake ecosystem subject to two steady states. *Ecological Economics* 59:429-439.
14. Kiseleva, T. & F.O.O. Wagener. 2008. Bifurcations of optimal vector fields in the shallow lake system. CeNDEF Working paper 08-01:1-7.
15. Kossioris, M. et al. 2008. Feedback nash equilibria for non-linear differential games in pollution control. *Journal of Economic Dynamics & Control* 32:1312–1331.
16. Dockner, E. & F.O.O. Wagener. 2008. Markov-perfect Nash equilibria in models with a single capital stock. CeNDEF Working paper 08-09:1-33.
17. Dechert, W.D. and S. O'Donnell. 2006. The stochastic lake game: A numerical solution. *Journal of Economic Dynamics & Control* 30:1569-1587.
18. Salerno, G., R. Beard & S. McDonald. 2007. The political economy of shallow lakes. Preprint.