HAL
open science

Extending the viability theory framework of resilience to uncertain dynamics, and application to lake eutrophication

C. Rougé, Jean-Denis Mathias, G. Deffuant

To cite this version:

C. Rougé, Jean-Denis Mathias, G. Deffuant. Extending the viability theory framework of resilience to uncertain dynamics, and application to lake eutrophication. Ecological Indicators, 2013, 29, p. 420 - p. 433. 10.1016/j.ecolind.2012.12.032 . hal-00830714

HAL Id: hal-00830714

https://hal.science/hal-00830714v1

Submitted on 5 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Author-produced version of the paper published in Ecological Indicators, vol.29, June 2013, pages 420-433
Original publication available at http://www.sciencedirect.com (Elsevier), DOI : http://dx.doi.org/10.1016/j.ecolind.2012.12.032

Extending the viability theory framework of resilience to uncertain dynamics, and application to lake eutrophication

Charles Rougé[1], Jean-Denis Mathias[1] and Guillaume Deffuant[1]

December 2012

Abstract

Resilience, the capacity for a system to recover from a perturbation so as to keep its properties and functions, is of growing concern to a wide range of environmental systems. The challenge is often to render this concept operational without betraying it, nor diluting its content. The focus here is on building on the viability theory framework of resilience to extend it to discrete-time stochastic dynamical systems. The viability framework describes properties of the system as a subset of its state space. This property is resilient to a perturbation if it can be recovered and kept by the system after a perturbation: its trajectory can come back and stay in the subset. This is shown to reflect a general definition of resilience. With stochastic dynamics, the stochastic viability kernel describes the robust states, in which the system has a high probability of staying in the subset for a long time. Then, probability of resilience is defined as the maximal probability that the system reaches a robust state within a time horizon. Management strategies that maximize the probability of resilience can be found through dynamic programming. It is then possible to compute a range of statistics on the time for restoring the property. The approach is illustrated on the example of lake eutrophication and shown to foster the use of different indicators that are adapted to distinct situations. Its relevance for the management of ecological systems is also discussed.
Keywords: Resilience; Viability theory; Stochastic dynamics; Probability of resilience; Dynamic programming

1 Introduction

Resilience has growingly been regarded as a central concept for many ecological systems, as well as for many human systems relying on ecosystem services. Indeed, resilience is related to the continued existence and sustainability of these systems in an era of upcoming widespread changes [e.g. 15, 30]. Since the original definition of resilience in ecology by [21], there has been a flurry of definitions of the concept in related fields [5], while it can be tackled at different levels of abstraction [6]. In this context, and in order not to lose what is meant by resilience, one central challenge is to introduce generic computational frameworks that can at the same time accurately reflect the concept and produce case-specific indicators.

[1]Irstea, UR LISC Laboratoire d'ingénierie des systèmes complexes, 24 avenue des Landais - BP 20085 Aubière, F-63172, France

This paper aims at tackling this issue for discrete-time stochastic dynamical systems. It builds on the viability theory framework of resilience, initially introduced in [26] for deterministic systems, and to which a book has recently been devoted [11]. It focuses on extending this framework to the stochastic case through the use of computational techniques based on dynamic programming. The viability framework generalizes the current mathematical definitions of the concept of resilience without betraying its intuitive sense. In this respect, we consider that the viability approach responds to the call from Brand and Jax [5] for a clear descriptive definition of the term of resilience, to avoid an eventual loss of its conceptual content as its use as a boundary object is spread around different research communities, for instance ecology and social sciences [1, 15, 2]. Besides, one major advantage of the viability approach is to provide with ways of computing relevant courses of action to ensure the resilience of the system. The extension presented here preserves this aspect.

The viability framework of resilience starts with the idea that it is essential to define the resilience of "what" to "what" [6]. It assumes that we are interested in the resilience of a property of the dynamical system, this property being called resilient to a perturbation if the system is able to restore it if it is lost, then keep it. In mathematical terms, a property can be described as a subset of the system state space, delimited by constraints. Resilience is related to the ability to come back and stay in this subset. Hence, this approach particularly emphasizes the possibility for restoring the property after losing it, which in our view is what distinguishes resilience from the concept of stability [17].

In this respect, the viability framework is distinct from the initial definition of resilience in ecology by [21] as the amount of perturbation a system can withstand while keeping its properties and functions. A direct implication of such a definition is the emphasis put on the considered properties and functions. The reason is that the considered resilience is of the system itself, hence of the properties and functions that define its identity. When considering the resilience of any property, possibly a minor one, we open the possibility for the system to lose the property for a while, and then to restore it. We claim that focusing on the aspects of recovery and restoration, which is put forward in the viability based framework, is pivotal to the resilience concept. There is indeed a variety of studies that relates resilience to measures of the recovery capacity of a system [e.g. 20, 25, 22, 23].

As explained in [11], the viability framework generalizes the mathematical attractor-based definitions of resilience [e.g. 4] which focus on regime shifts. Indeed, in these latter definitions, the property of interest is identified as a set of states located around selected attractors which represent a desirable regime. Resilience indicators are related to the size of the attractor basins of these attractors, namely the part of space where the system avoids to fall into bad attractor basins. If one defines the property of interest as a subset around the good attractors, then the viability framework provides similar results to the ones of the attractor based definition. However, the property of interest can be defined as any subset of the state space and therefore offers wider possibilities, especially in cases where no attractor would exist.

Moreover, the viability framework of resilience provides operational tools for computing policies of actions (or feedback rules) to keep or restore a property of interest, namely by driving the system back to the desirable subset and keeping it there. Alternative mathematical frameworks for computing resilience at best suppose that a policy of action is already defined. This is a very important practical advantage of the viability framework, as illustrated in several case studies in [11]. As any other

framework, viability has its limits and drawbacks, namely:

1) The current numerical approaches for computing the resilient states and the restoration policies are very demanding computationally because they require to discretize the state space (and also the action space). Therefore, in practice only the dynamical models with very few degrees of freedom are tractable (see [11] for further information).

2) The current viability framework only considers one-time perturbations in otherwise deterministic dynamical systems, thus not taking into account the potential impact of other uncertainty sources while computing resilience to these events.

This paper aims at overcoming the latter limit and proposes an extension of the viability framework of resilience to stochastic dynamical systems. It is based on the key assumption that uncertainty, whatever its source, can be mathematically described in two different ways. On the one hand, some of the uncertainty can be described by stochastic processes, which can be assessed under the form of probability distribution functions (pdf) and reduced by experience [3, 31]. Since this part of the uncertainty can be embedded into the dynamics, the system can be managed so as to be made robust to it. On the other hand, there is invariably a part of unpredictability that escapes such assessments [18, 29], as emphasized by the expression "uncertainty and surprise" found in the literature [15, 16, 2]. This second uncertainty source generally refers to events that initiate potentially major disturbances, and to which the resilience of the system must be assessed.

From a technical point of view, the present work uses the stochastic extensions to viability theory in the discrete-time case [8, 13] and builds on the robust and stochastic viability frameworks, both based on dynamic programming, and which have been successfully applied to fisheries management [12, 9, 14, 27]. One can expect three main differences between the deterministic and stochastic viability frameworks:

1) The viability kernel is a subset of the state space of a deterministic system in which the system can keep the property if it remains undisturbed. Stochastic viability kernels only guarantee that the system will keep the property with a given probability by a given time horizon, but this is also a guarantee that the system is robust to the uncertainty sources described by the stochasticity of its dynamics.

2) The single optimal trajectory for reaching the deterministic viability kernel is replaced by a set of trajectories. This fosters the definition a probability of resilience, the probability of reaching the stochastic viability kernel by a given time horizon. Besides, the single resilience indicator that prevails in the deterministic case (e.g. the inverse of the restoration cost as in [26]) is replaced by a pdf for the times or costs of restoration. Resilience indicators are derived from statistics on this pdf.

3) Contrary to what happens with deterministic viability, the time horizon considered for the management of the system becomes of paramount importance for policy design. Policies become specifically designed to maximize the probability of being resilient by the time horizon, but are not in general meant to maximize this probability at any shorter horizon.

The paper is organized as follows. First, we recall the main concepts of the viability-based framework of resilience in the deterministic case, and illustrate it on a simple lake eutrophication model (Section 2). In section 3, we introduce the extension of the viability based framework of resilience to stochastic dynamics, defining the concepts of probability of resilience and different measures of restoration time and cost. Section 4 illustrates the approaches on the lake model, first for a given control strategy, then when controls are being optimized as resilience is computed. The relevance for the study of ecological systems is discussed in Section 5, before conclusions are drawn in Section 6.

2 Background: resilience in the deterministic case

The general problem of achieving resilience in limited time after a perturbations, considering constrained dynamics and controlled strategies, is laid out in discrete time in the deterministic case. The viability framework of resilience from [26] and [11] is then put into perspective.

2.1 Problem statement

2.1.1 Controlled dynamics

In the viability framework for resilience presented by [26], an important innovation is to introduce controls to explicitly account for the possibility to act on the system. In this framework, the policy is not fixed beforehand. Instead, the goal is to find policies that will make the system resilient. In discrete time, this means that at each time step, there is a set of possible actions that one must choose from, and a known transition equation between two consecutive dates. Let us note X the state space and x_t the state of the system at date t. Noting also $U(x_t, t)$ the set of available controls and u_t the chosen control value, a typical controlled discrete-time dynamical system can be written as:

$$x_{t+1} = g\left(x_t, u_t\right) \tag{1}$$

2.1.2 Resilience to "what"?

In this framework we focus on resilience to a given perturbation. There is no assumption on the nature or amplitude of the perturbation. Indeed, we place ourselves in the post-perturbation state of the system. Resilience of any state to any perturbation can then be assessed by looking at the new state of the system after the perturbation.

2.1.3 Management objectives and constraints

The system has some properties that are deemed desirable. We assume that these properties can be mathematically translated into state constraints, which define a set of desirable states noted K. The general goal then becomes to control the system so it stays within K.

After a perturbation, the goal becomes to ensure that the system gets back to a state where its dynamics are going to keep it within K for as long as it is not disturbed again. Management also has a time frame, or time scale of interest, which we will note T. This is the time by which the system's properties ought to be restored.

2.1.4 Management strategies

The objectives are achieved through management strategies. A strategy can be represented by a function which associates a control to any date $0 \leq t \leq T-1$ and to any state x. A strategy can be described at each time step by a feedback map, an application from X into $U(x,t)$ (or from K if we only want to see whether the system leaves K). In this work, strategies will be determined depending on the time horizon. Consequently, we choose to reference a feedback maps are referenced by their time distance to the horizon. A strategy is given, in chronological order, by the following succession of feedback maps:

$$f = (f_T, f_{T-1}, \ldots, f_1) \tag{2}$$

Controls at any date $t \leq T$ are deduced from equation (3):

$$\begin{cases} u_t = f_{T-t}(x) \\ u_t \in U(x,t) \end{cases} \tag{3}$$

We can also introduce the notation $F(T)$ for the set of all strategies f with a time horizon T. For a given f, it is possible to recursively compute all the states from the initial state x_0 to the final state x_T. Indeed, equation (1) can be written anew using equation (3) through:

$$x_{t+1} = g(x_t, f_{T-t}(x_t)) \tag{4}$$

Thus, we can define the trajectory g_f starting from the initial state x_0 and using the strategy f as:

$$\forall t \in [0,T],\ x_t = g_f(t, x_0) \tag{5}$$

Figure 1 summarizes the notations introduced in this section for a horizon T and the associated strategy f. The choice of a strategy (equation (2)) allows for the computation of controls (equation (3)). Successive states can then be computed through equation (1), leading to the computation of trajectories (equation (5)).

2.2 The viability kernel

Prior to assessing resilience, one must determine which are the desirable states of K for which the dynamics can keep the system properties until another perturbation occurs, however long it may take for that to happen. In order to guarantee that, the properties need to be kept for a length of time much greater than the time scale of interest, T. Resilience will then be the ability to reach one of these states.

Under the framework of viability theory, the set of all the states for which there is a control strategy such that the system can be maintained inside the set of desirable states throughout a period of time $\tau >> T$ is called the viability kernel. In discrete time, it can be formally defined as the set of initial states for which there exists a trajectory that does not leave K during τ time steps:

$$\text{Viab}(\tau) = \{\ x_0 \in K \mid \exists f \in F(\tau),\ \forall t \leq \tau,\ x_t = g_f(t, x_0) \in K\} \tag{6}$$

and for simplicity, the notation Viab will be used instead of Viab(τ) in the remainder of Section 2.

In practice, several algorithms exist to determine which states belong to the viability kernel [28, 10]. Its computation also yields the set of controls which maintain the system, which are called the viable controls. Thus, it incorporates the impacts of management policies, and implicitly optimizes them.

2.3 Viability-based definition of resilience

This section puts into perspective the results from [26] and [11] about resilience to a single perturbation in an otherwise deterministic system. Resilience is associated to the possibility of getting back to the viability kernel in a relevant time frame, denoted by the horizon T.

Let us first consider a given strategy f at horizon T, with the only constraint that at any date and for any state within the viability kernel $f_{T-t}(x)$ is a viable control. One can then define what is called a resilience basin in [11]: the set of states for which the system is brought back to the viability kernel in a horizon $t \leq T$. This notion can first be written for a given strategy:

$$B^{\text{res}}(f,t) = \{\ x \in X \mid g_f(t,x) \in \text{Viab}\ \} \tag{7}$$

For $t_1 < t_2$, since any trajectories that reaches the viability kernel by t_1 can also reach it if given a greater time horizon, we have $B^{\text{res}}(f,t_1) \subset B^{\text{res}}(f,t_2)$. The union of all the resilience basins is $B^{\text{res}}(f,T)$, and it is the set of resilient states for the strategy f, from which the system can recover by getting back to the viability kernel in T time steps or less.

Yet, one has many strategies at her disposal, and may want to find the set of all the states for which it is possible to bring the system back to the viability kernel in a horizon $t \leq T$. This brings about a different definition of a resilience basin:

$$B^{\text{res}}(t) = \{\ x \in \mathbb{R}^n \mid \exists f \in F(t),\ g_f(t,x) \in \text{Viab}\} \tag{8}$$

The implicit difference between equations (7) and (8) is that in the latter, controls have been optimized to find trajectories that ensure the resilience of the system by the chosen time horizon. In this work, we choose to use dynamic programming to carry out this optimization. Dynamic programming is a recursive algorithm which enables to optimize a value function, noted $V_d(t,x)$ in this deterministic case, at each horizon t. It progresses backwards from date T (horizon 0) to the initial date (horizon T). It yields the same results in the deterministic case as algorithms like KAVIAR [10] used by [11], but will be extended to the uncertain case in Section 3. Like any recursive algorithm, it works based on an initial equation and a transition equation. Here, initialization comes from the fact that resilience at a time horizon of 0 is the fact of belonging to the viability kernel:

$$V_d(0,x) = \begin{cases} 1 & \text{if} \quad x \in \text{Viab} \\ 0 & \text{if} \quad x \notin \text{Viab} \end{cases} \tag{9}$$

Then the algorithm progresses recursively by incrementing the horizon considered from 0 to T thanks

to the following transition equation:

$$\forall t \in [1, T],\ V_d(t, x) = \max_{u \in U(x, T-t)} V_d(t-1, g(x, u)) \tag{10}$$

This value function can only take values 0 and 1. The work by [13] links the computation of V_d with the resilience basins $B^{\text{res}}(t)$ (see Appendix A for the proof):

$$B^{\text{res}}(t) = \{\ x \in X \mid V_d(t, x) = 1\} \tag{11}$$

The set of resilient states is the resilience basin $B^{\text{res}}(T)$. The definitions introduced in this section are graphically summarized in Figure 2. A great strength of the viability framework is to dynamically compute the optimal controls and the resilience basins at the same time. One can thus the restoration time at $x \in B^{\text{res}}(T)$, which can be handily defined as the minimal horizon for which a trajectory starting at x can reach the viability kernel:

$$t^*(x) = \min_{t \leq T} \left\{ t | x \in B^{\text{res}}(t) \right\} \tag{12}$$

and the associated optimal policies are defined by recurrence. For x outside the viability kernel, and $t^*(x) > 0$, any control that enables the system to reach $B^{\text{res}}(t^*(x) - 1)$ preserves the possibility for the properties to be restored by the horizon t^*. Besides, the controls only need to be set once, so that in the deterministic case, the same feedback map can be used regardless of the horizon. This will not be true in the stochastic case which we introduce now.

3 Resilience computations in uncertain discrete-time systems

This section presents the extension of the previous framework to stochastic systems, which is the main contribution of this paper. Uncertainty is first introduced into the modeling framework described by Section 2.1. A stochastic equivalent of the viability kernel is introduced to describe the safe states of the system. Since one can no longer guarantee that a set can be reached with probability one after a minimal amount of time, we choose to focus on maximizing the probability of being resilient by reaching the stochastic viability kernel. We thus introduce the notion of probability of resilience, before showing how quantities that can be related to a variety of possible resilience indicators can be handily computed thanks to this notion.

3.1 Incorporating uncertainty

Uncertainty in a modeling framework such as that of Section 2.1 can arise from numerous sources [31]. Yet, here we are concerned with how uncertainty can be modeled, rather than where it comes from. Knowledge about uncertainty can be imperfect [29], so that no pdf can capture all of the uncertainty, especially that related to large events. As mentioned in the introduction, the key assumption used here is to consider that aside from large uncertain perturbations which might happen from to time to time, the rest of the uncertainty can be modeled at each time step using random variables for which a pdf can be defined. Then, we compute resilience to a given large, unknown event while taking into

account the sequential effects of smaller events which may be described in stochastic terms at each time step.

When considering uncertainty, the dynamic g of equation (1) has to be modified. Following [13], let us assume the existence of a known dynamic g which incorporates a vector ε of all uncertainties, whether they concern the dynamic itself or the state, control or parameter variables. We get:

$$x_{t+1} = g(x_t, u_t, \varepsilon_t) \tag{13}$$

Let us now suppose that the state space X has been discretized, which is the case in practice. Discretization methods are not within the scope of this paper. We note $\hat{X}$ the points of the discrete grid.

The rest of Section 2.1 is kept unchanged. In particular, the notion of strategy $f \in F(T)$ is kept; unlike what happens in deterministic viability, one should not expect the feedback map to be the same for two different values of the horizon T. We also keep the notion of trajectory associated to a strategy, but we are now faced with a set of trajectories which depends on the events ε_0, ε_1, etc... Thus, $g_f(t, x_0)$ becomes a random variable.

The use of dynamic programming as an essential tool of stochastic viability theory [13] imposes the pdfs ε_t to be uncorrelated with one another. This is because dynamic programming essentially requires to travel backwards in time. What is more, in this framework, feedback optimization is only guaranteed when the ε_t are independent and identically distributed (i.i.d.), which we are going to assume from now on. Uncertainty can then be made implicit in equation (13):

$$x_{t+1} = g_\varepsilon(x_t, u_t) \tag{14}$$

For a given state x_t and decision u_t, the probability of the state value being y at date $t+1$ will be noted using the discrete probability $\mathbb{P}(g_\varepsilon(x_t, u_t) = y)$. Stochastic viability kernels will now be introduced to describe robustness to the uncertainty ε.

3.2 Robustness to uncertainty

With the introduction of uncertainty, it is not possible in general to find states that ensure with unit probability that a system will retain its properties if there is no major perturbation before a date τ. Instead, one can introduce the stochastic viability kernel [8], defined as the set of initial states for which the system has a probability β or higher of keeping its properties for a duration τ. Noted $\mathrm{Viab}(\beta, \tau)$, it can be formally defined by the following equation:

$$\mathrm{Viab}(\beta, \tau) = \{\ x_0 \in K \mid \exists f \in F(\tau),\ \mathbb{P}\left(\forall t \in [0, \tau], x_t = g_f(t, x_0) \in K\right) \geq \beta\ \} \tag{15}$$

For instance, $\mathrm{Viab}(0.99, 100)$ is the set of initial states such that the system has at least 99% chance of avoiding the loss of its properties for at least a hundred time steps. An interesting case arises when $\beta = 1$: indeed $\mathrm{Viab}(1, \tau)$ is the set of initial states x_0 for which the system is kept in a desirable state with certainty during τ time steps. This set is called the *robust viability kernel* and is an analog of the viability kernel defined in Section 2.2. The stochastic viability kernel is computed thanks to the

dynamic programming algorithm proposed by [13] (see A for details).

Like for the deterministic viability kernel, the feedbacks are being optimized, but optimization now occurs at each time horizon on the interval $[0, \tau - 1]$. Thus, dynamically computing the optimal management strategy while looking for the robust states of the system remains a major advantage of the viability approach.

Sets like $\mathrm{Viab}(\beta, \tau)$ with β close or equal to 1 and $\tau \geq 100$ can be seen as the set of system states where the desirable properties of the system hold (almost) certainly despite the onset of quantifiable uncertainties or disturbances. Therefore stochastic viability kernels are sets where the system properties are robust to the uncertainty sources modeled by the process ε, and the parameters β and τ describe the extent of such robustness. The term ε-robustness is here used to designate the property of state within a stochastic viability kernel. However, the focus of this work is to compute resilience to a large unexpected perturbation, and a new notion is needed to extend to the stochastic case the viability framework, which precisely dealt with resilience to such events in the deterministic case.

3.3 Probability of resilience

Let us now assume that one has computed a stochastic (or robust) viability kernel $\mathrm{Viab}(\beta, \tau)$, the set of states for which the system's properties are guaranteed to be ε-robust. The conceptual notion of resilience can then be operationalized through the possibility of reaching this set in a given time frame T. This is very similar to what is done in the deterministic case, except that reaching $\mathrm{Viab}(\beta, \tau)$ now guarantees ε-robustness to quantifiable uncertainty. Besides, the considered horizon T has to be small compared to the time scale τ during which the system can be maintained with a high probability, so that the condition $T << \tau$ holds, like in Section 2.

For a given strategy $f \in F(T)$, the probability of getting into a given $\mathrm{Viab}(\beta, \tau)$ by a date $t \leq T$ is the capacity to recover from a perturbation under this strategy. It is noted $\mathbb{P}_{\mathrm{Res}}(f, t, x)$:

$$\mathbb{P}_{\mathrm{Res}}(f, t, x) = \mathbb{P}\left(\exists j \in [0, t], g_f(j, x) \in \mathrm{Viab}(\beta, \tau)\right) \tag{16}$$

and from equation (2), the controls that are applied between dates 1 and t are those of the successive feedback maps $(f_T, f_{T-1}, \ldots, f_{T-t})$.

Like in the deterministic case, the objective of viability is to find the feedbacks that maximize the probability of resilience, but now specifically as the probability of reaching $\mathrm{Viab}(\beta, \tau)$ by date T. This probability of resilience, noted $\mathbb{P}_{\mathrm{Res}}(T, x)$, is thus given by the value of $\mathbb{P}_{\mathrm{Res}}(f, T, x)$ when the strategy is optimal:

$$\begin{aligned} \mathbb{P}_{\mathrm{Res}}(T, x) &= \max_{f \in F(T)} \mathbb{P}_{\mathrm{Res}}(f, T, x) \\ &= \mathbb{P}_{\mathrm{Res}}(f^*, T, x) \end{aligned} \tag{17}$$

where the optimal strategy f^* is not necessarily unique. The probability of resilience $\mathbb{P}_{\mathrm{Res}}(T, x)$ represents the probability of being resilient within the considered time frame. If $\mathbb{P}_{\mathrm{Res}}(T, x) = 1$, the system can always be made ε-robust, and it is resilient.

To compute the probability of resilience in this stochastic case, we use a value function V_s depending

on the horizon t and state x, which is the exact analog of the deterministic V_d:

$$\begin{cases} V_s(0,x) = \begin{cases} 1 & \text{if} \quad x \in \text{Viab}(\beta,\tau) \\ 0 & \text{if} \quad x \notin \text{Viab}(\beta,\tau) \end{cases} \\ \forall t \in [1,T],\, V_s(t,x) = \begin{cases} 1 & \text{if} \quad x \in \text{Viab}(\beta,\tau) \\ \max\limits_{f_t(x) \in U(x,T-t)} \left(\sum\limits_{y \in \hat{X}} \mathbb{P}(g_\varepsilon(x, f_t(x)) = y)\; V_s(t-1,y) \right) & \text{if} \quad x \notin \text{Viab}(\beta,\tau) \end{cases} \end{cases} \tag{18}$$

and the value function V_s yields the maximal value of the probability of resilience (proof to be found in A):

$$\mathbb{P}_{\text{Res}}(T,x) = V_s(T,x) \tag{19}$$

and the associated optimal strategy f^* is given by the feedback maps $f^*_{T-t} : x \to u_t$, computed through equation (18) at each time step. This strategy $f^* = (f^*_T, f^*_{T-1}, \ldots, f^*_1)$ yields the maximal probability for the system to be resilient at the horizon T. It can then used to compute the probabilities $\mathbb{P}_{\text{Res}}(f^*, t, x)$ of recovering the ε-robustness of the system's properties by a date $t \leq T$.

The feedback one has to apply to a given state at a given date depend on the distance to the horizon T, and thus, on the horizon itself. Indeed, the probabilities $\mathbb{P}_{\text{Res}}(f^*, t, x)$ are not necessarily the maximal probability of reaching the stochastic viability kernel by t. For that, one should use the optimized feedbacks $(f^*_t, f^*_{t-1}, \ldots, f^*_1)$ for the respective dates $(0, 1, \ldots, t)$. These feedbacks would yield $\mathbb{P}_{\text{Res}}(t, x)$ which the maximal probability of being resilient with a time horizon t. Yet, one should keep in mind that between the initial date and t, the strategies $(f^*_t, f^*_{t-1}, \ldots, f^*_1)$ and $(f^*_T, f^*_{T-1}, \ldots, f^*_{T+1-t})$ are in general different, and that only the latter maximizes the probability of being resilient at date t.

The definitions introduced in this section are graphically summarized in Figure 3. One can notice differences with the resilience basins one can compute in the deterministic case (Section 2.3 and Figure 2). Due to the many possible trajectories for the optimal feedback strategy, resilience is not in general a certain property any more. As a consequence, minimizing the time to get back to the viability kernel does not make obvious sense like it does in the deterministic case. Instead, we choose to maximize the probability of getting back at a given horizon instead.

Besides, the probability of resilience is highly dependent on the prior choice of the β and τ. Indeed, the larger β and τ, the smaller the stochastic viability kernel, and the longer it takes to reach it. Yet, the more draconian the robustness criteria, the better the robustness of the system when it reaches that set, and the more meaningful it is to see resilience as the probability of reaching it. Thus when setting the parameters β and τ, one has to consider a tradeoff between the quality of the resilience and robustness of the properties of the system, and how costly these are to achieve. This means resilience to a large, unpredictable perturbation is defined in tandem with ε-robustness.

Eventually, the optimal feedbacks used to compute the ε-robust states and the resilient ones are not the same in general. They correspond to two distinct problems: the first is to guarantee the properties of the system against "usual" uncertainty, while the second is to make the system robust again following a perturbation. In both cases, the feedbacks are dependent on the time horizon considered.

3.4 Resilience-related indicators

Computed for a given T, the quantities $\mathbb{P}_{\text{Res}}(f^*, t, x)$ form a set of indicators that gives the cumulative distribution function of the time needed to get to a set where the desirable properties of the system are guaranteed to a comfortable extent against uncertainty. Yet, measures of resilience are often given through a single performance indicator. Thus, while they represent well what resilience is at the conceptual level, these quantities are not necessarily relevant operational indicators. However, they are a basis for computing virtually any indicator that represents resilience based on the time taken to reach a situation that keeps the properties of interest of a system. Indeed, the pdf of the time taken to reach an ε-robust set $\text{Viab}(\beta, \tau)$ is given by the difference $\left(\mathbb{P}_{\text{Res}}(f^*, t, x) - \mathbb{P}_{\text{Res}}(f^*, t-1, x)\right)$, which is the probability to reach this set exactly at date t.

Index computation depends on whether or not the loss of ε-robustness is reversible with unit probability by the date T. If there are states such that $\mathbb{P}_{\text{Res}}(T, x) < 1$, one can directly relate this quantity to the resilience of the system. Yet further relevant indicators related to resilience exist, especially (but not only) when $\mathbb{P}_{\text{Res}}(T, x) = 1$. One such indicator can be the expected value of the time needed to acquire the resilience property using the strategy f^*:

$$E(t|f^*, x) = \sum_{t=0}^{T} (t+1) \left(\mathbb{P}_{\text{Res}}(f^*, t+1, x) - \mathbb{P}_{\text{Res}}(f^*, t, x)\right) \tag{20}$$

It should be noted that f^* is not necessarily the strategy that optimizes the expected entry time, and it is no equivalent to the entry time defined in the deterministic case by equation (12). Such an equivalent could be provided by tailoring the considered horizon to the state of the system, still keeping in mind that it means the feedback maps used will then change. For x, one would pick the shortest horizon t_α such that the probability of resilience $\mathbb{P}_{\text{Res}}(t, x)$ guarantees that the ε-robustness of the system at the horizon t_α with a confidence level α. When it exists, t_α can be defined by:

$$t_\alpha(x) = \min_{0 \leq t \leq T} \left\{t | \mathbb{P}_{\text{Res}}(t, x) \geq \alpha\right\} \tag{21}$$

and because of equation (19), it can be computed using yet again the value function V_s:

$$t_\alpha(x) = \min_{0 \leq t \leq T} \left\{t | V_s(t, x) \geq \alpha\right\} \tag{22}$$

Both indicators $E(t|f^*, x)$ and $t_\alpha(x)$ can be used in association with the notion of resilience, and a resilience indicator can even be defined as a decreasing function of one of them. Then, this decreasing function is often the inverse function [20, 26]. In a more general way, when resilience can be achieved with a nonzero probability, a useful indicator is the cost incurred while driving the system back to the viability kernel. One can choose whether to associate it to the notion of resilience or to define resilience as a decreasing function of this cost much like in [26] (e.g. its inverse).

The equations proposed here provide one with various ways to compute resilience and related indicators while explicitly incorporating the potential effects of uncertainty. They are also an illustration of the power of dynamic programming tools when it comes to carrying out resilience computations.

4 Application to an uncertain lake model

The lake eutrophication problem can be tackled through simple nonlinear dynamics, and has been well-studied [7, 6, 24]. Here we choose, for illustrative purposes, to use the deterministic model from [26] and add uncertainty to it. Resilience computations introduced by [26] and [11] are outlined, and the added value of considering uncertainty is highlighted. The dynamics and resilience problem are laid out first, before being solved for a given strategy. Then, the interest of using stochastic viability to optimize the control strategy is emphasized.

4.1 The problem of lake eutrophication

Using the publications cited above, this section introduces the case of lake eutrophication under its simplest formulation, that is, not considering the mud dynamics. The problem is formulated much like in Section 2.1, but introduces uncertainty while introducing the system's dynamic.

4.1.1 Uncertain controlled dynamics

The state variables are the phosphorus concentration p in the lake, and the phosphorus input rate l. In continuous time θ, the evolution of phosphorus concentration reads as follows [26]:

$$\frac{dp}{d\theta} = -bp + l + r\frac{p^q}{p^q + m^q} \tag{23}$$

where the parameter b is the phosphorus sink rate (e.g. the quantity that flows out of the lake), r is the maximal recycling rate by algae, m is the value of p for which the recycling term $(rp^q)/(p^q + m^q)$ is half its maximal value, and q is a dimensionless parameter, set at $q = 8$ in several studies [e.g. 26, 19].

Through dimensional analysis, the parameters b and m can be eliminated while p, l and r are turned into their respective dimensionless equivalents: the state variables P and L, and the only remaining parameter is R (see Appendix B for details). θ is also replaced by the dimensionless time t, and we get the new continuous-time equation:

$$\frac{dP}{dt} = -P + L + R\frac{P^8}{P^8 + 1} \tag{24}$$

Let us now tackle uncertainty. For the sake of simplicity, only phosphorus input uncertainty will be considered in this example. It represents, for instance, the uncertainty that is due to the soil storage of phosphorus. A discrete-time decomposition (for which the continuous-time equivalent is given in C.1) between the mean input rate L^* and a deviation is introduced:

$$L_t = L_t^* + \sigma\varepsilon_t \tag{25}$$

Here ε is considered i.i.d., and it is modeled by a standard normal distribution $\mathcal{N}(0, 1)$, due to the common use of such distributions. σ represents the standard deviation of the noise. Since the phosphorus input L is not updated by the dynamic of equation (23), choosing L^* instead of L as a state variable enables the modeler to project the uncertainty over one dimension (P) rather than two (L

and P). It therefore reduces the computational complexity of the problem. From now on, the state x will be the couple (L^*, P).

The mean phosphorus input L^* can be modified due to modifications in farmers' behavior, changes in agricultural technology, a combination of both, or other factors. Such modifications take time to take full effect, so that the rate of change in L^* is bounded. One can write:

$$\begin{cases} L^*_{t+1} = L^*_t + u_t \\ u_t \in U(x_t) = [U_{\min}(x_t), U_{\max}(x_t)] \end{cases} \tag{26}$$

where for any x_t, $U_{\min}(x_t) \leq 0$ and $U_{\max}(x_t) \geq 0$. In this application, the set of possible controls depends only on the state, and not on time.

The discrete time equation for the evolution of the lake is, based on the derivations from Appendix C.3:

$$P_{t+1} = g(P_t, L^*_t, u_t, \varepsilon_t) \tag{27}$$

We discretize the state space with a resolution $\Delta P = \Delta L^* = 0.01$, so that for a given value of the control it is possible to directly compute the probability for the state to be closest to any given point of the grid at the end of the time step. One can thus describe the evolution of the lake state from any point $x_t = (L^*_t, P_t)$ of the discrete grid to another.

4.1.2 Resilience to "what"?

By considering the post-perturbation state, this framework can allow for computing the resilience to any given perturbation. In practice, a perturbation is any event not taken into account by equation (25). Later in this work, the focus will be on sudden, single increases in the value of L^* or P. The former represents an increase in the mean quantity to reach the lake, which can be due to the existence of a massive new phosphorus source, and its amplitude will be noted L^*_{per}. The latter rather represents a one-time arrival of phosphorus into the lake, which triggers an immediate increase in the phosphorus concentration P, and its amplitude will be noted P_{per}.

4.1.3 Management objectives and constraints

Phosphorus inputs are the by-product of the use of fertilizers, which benefit to farmers through an improved land productivity. However, a lake can have two regimes, and phosphorus concentration has been found to trigger a regime shift. Namely, the switch is from the oligotrophic or clear water regime, in which both ecologic and economic benefits from the lake are high, to the eutrophic or turbid water regime in which algae blooms cause oxygen depletion, leading in turn to a so-called dead lake. Therefore, this is an undesirable regime shift, and we can write the set of desirable states, which we note K, as:

$$K = [L^*_{\min}, L^*_{\max}] \times [0, P_{\max}] \tag{28}$$

where $L^*_{\min}$ corresponds to the minimum quantity of phosphorus needed for farming to remain economically viable; $L^*_{\max}$ is the maximum amount of phosphorus that farmers are willing to use; and

$P_{\max}$ is the threshold above which the lake turns eutrophic.

For the rest of this work, bounds for L^* are set at $L^*_{\min} = 0.1$ and $L^*_{\max} = 1$ like in [26]. We assume that the farmers will not accept to modify the value of L^* beyond the prescribed bounds, so that we can define $U_{\min}(x)$ and $U_{\max}(x)$ depending on those bounds and the maximal amplitude M at which L^* can be modified:

$$U_{\min}(x) = \max\{-M, L^*_{\min} - L^*\} \tag{29}$$

$$U_{\max}(x) = \min\{M, L^*_{\max} - L^*\} \tag{30}$$

Besides,the recycling term $(RP^q)/(P^q + 1)$ increases sharply around $P = 1$, and this change in recycling rate characterizes the transition from the oligotrophic to the eutrophic regime, so that 1 is a plausible value for $P_{\max}$ in equation (28).

The goal of management is to keep the lake in a state in which it can be maintained in K for a long time, or alternatively, to reach such a state if a perturbation were to occur. Management considers a time horizon T to bring the system back to a state where it can be maintained for a long time. The certainty and uncontrolled case are respectively given by $\sigma = 0$ and by $M = 0$.

4.1.4 Management strategies

To achieve the management objectives, let us introduce management strategies as defined by equations (2) through (5). We note f the considered strategy, and T the time horizon at which we will consider managing the system:

$$\begin{cases} u_t = f_{T-t}(P_t, L^*_t) \\ u_t \in U(x) \end{cases} \tag{31}$$

In practice, one can expect strategies to mainly aim at reducing the mean phosphorus input L^*. This can be seen using the position of the equilibria of equation (27) in the classical certain and uncontrolled case ($\sigma = 0$, and $M = 0$, so that $L = L^*$) (Figure 4). Depending on the value of the parameter R, there can be three types of behavior for the lake. These have been well studied in the literature. When R is low, for instance $R = 0.4$, there is only one value of P which is an equilibrium for each value of L, and the lake is called reversible. For higher values, such as $R = 1.2$, there can be alternate state states for the same value of L, making a regime shift much more difficult to reverse. The lake is then called hysteretic. Finally when R is even higher, for instance at 2, the lake can be called irreversible, because any transition towards the eutrophic regime cannot be reversed. The reversible case is only weakly nonlinear, so that the remainder of this paper will focus exclusively on the hysteretic and irreversible cases, with the same respective values of R (1.2 and 2) as in Figure 4.

4.2 Resilience for a given strategy

In this section, resilience is explored using predefined strategy. The focus will be on 1) strategy $f^{(1)} \equiv 0$ where no management strategy is implemented, and 2) a purely reactive strategy $f^{(2)}$ where an effort is made once the lake has become eutrophic, by reducing the mean phosphorus input as long as the

constraint $L^* \geq L^*_{\min}$ is respected:

$$\forall t \in [0,T],\ f^{(2)}_{T-t}(x) = \begin{cases} 0 & \text{if} \quad x \in K \\ U_{\min}(x) & \text{if} \quad x \notin K \end{cases} \tag{32}$$

where we set $M = 0.05$. In this paragraph, we also set $R = 1.2$. We will first explore the advantage of using viability kernels rather than attractors to describe the states that are robust to uncertainty (still using the term of ε-uncertainty), then we will give resilience indicators as introduced in Section 3 for these two strategies.

4.2.1 Stochastic viability kernels and their advantages

The search for ε-robust states considers the trajectory of the system only as long as it does not exit the set K of desirable states. Since both the strategies considered here yield $f \equiv 0$ within K, we only need to consider strategy $f^{(1)} \equiv 0$, and $f^{(2)}$ will give exactly the same results.

Figure 5 shows how a system set at a stable desirable attractor of the (L, P) plane can switch towards an undesirable state in a short time span when uncertainty is present. That had already been noted by [19]. Thus, uncertainty causes a description of the dynamics of the system using stable equilibria and their basins of attraction to become precarious. Switching variables to (L^*, P) does not change this fundamental fact, even though they do allow for the projection of uncertainty over a single state variable instead of two (Figure 6). Indeed, while uncertainty does not influence the value of L^*, it does interfere strongly with the position of the equilibria in the (L^*, P) plane. For a system in the same initial state as in Figure 5, the 95% confidence interval (CI) associated with its state after 5 time steps (Figure 6) shows that the oligotrophic attractor only describes the dynamic very poorly. Using equilibria that change at every time step is very unpractical, and moreover, one can perceive that the system is very unlikely to be on a stable equilibrium at any date.

Since it situates a system with respect to a set rather than with respect to a point, viability theory is well-suited to describe out-of-equilibrium situations [26, 11]. One can easily represent different degrees of ε-robustness of the system by computing stochastic viability kernels $\text{Viab}(\beta, \tau)$ for different values of β and τ (Figure 7). Equilibria for the mean case $\varepsilon_t = 0$ are still shown for an easier understanding. The limit case $\sigma = 0$ is also represented: it corresponds to the deterministic viability kernel.

As expected, for a given value of β and σ, the higher τ, the smaller the size of the kernel. Likewise, for (β, τ), the higher the scale σ of the uncertainty, the smaller the size of $\text{Viab}(\beta, \tau)$. The same result holds when β increases for given values of β and σ. Besides, the size decrease caused by an increase in τ is magnified by higher values of σ. Figure 7 shows that for a hysteretic lake, one can get nonempty stochastic viability kernels $\text{Viab}(\beta, \tau)$ with β close to one, e.g. $\beta = 0.99$, and a time horizon much more important than that of decision-making, e.g. $\tau = 10^3$. The set $\text{Viab}(0.99, 10^3)$ will be used from now on to describe the ε-robust states. We also set $\sigma = 0.1$.

4.2.2 Resilience

Let us now assume a management horizon $T = 30$ ($T << 10^3$) and look at the possibility to reach a ε-robust state after a perturbation with either strategy $f^{(1)}$ or $f^{(2)}$. We solely look at the post-

perturbation state of the system. The null strategy $f^{(1)}$ lets L^* stay constant over time, so that reaching $\text{Viab}(0.99, 10^3)$ is only possible if for a given L^*, there are values of P which are ε-robust. Hence, outside the stochastic viability kernel, resilience is only possible for a small portion of the state space (Figure 8).

For strategy $f^{(2)}$ however, there is a possibility to control the system when the constraints are violated, and a possibility for the system to get to a ε-robust state once L^* gets very low. Hence, with $M = 0.05$, most post-perturbations states are resilient with unit probability at the horizon $T = 30$ (Figure 9), highlighting the importance of management actions. The only exception concerns states which are not ε-robust if we consider $\beta = 0.99$ and $\tau = 10^3$, but for which the probability of a switch towards a eutrophic lake is relatively slim. Strategy $f^{(2)}$ then leaves the system in a precarious state until that switch happens. One can intuitively sense that in such cases, a strategy that would take action and reduce L^* before the lake becomes eutrophic would be more environmentally efficient. In fact, as described in Section 3, viability theory allows for searching the best strategy available.

4.3 Resilience with an optimal management strategy

Let us now fully exploit a major advantage of viability theory: the possibility to select the policy choices that maximize the probability of resilience at the horizon T. Through this whole section, we will have $\sigma = 0.1$ and $M = 0.05$, and the lake will be considered hysteretic $R = 1.2$ unless mentioned otherwise.

4.3.1 Implications for stochastic viability kernels

Thanks to the work by [13], feedbacks inside K can be computed by dynamic programming so as to maximize the probability of keeping the lake in a desirable state. Whatever the value of τ and for all three values of R considered, the method yields the same optimal feedback map at all dates, and it is noted f^*:

$$f^*(x) = U_{\min}(x) \tag{33}$$

Yet again, we choose to describe ε-robustness using $(\beta, \tau) = (0.99, 10^3)$. As expected, the associated viability kernel is much bigger than with strategies $f^{(1)}$ or $f^{(2)}$, as evidenced on Figure 10. It in fact encompasses states that we deemed "precarious" when using strategy $f^{(2)}$, so that the optimal strategy makes them "robust" to measurable uncertainty.

4.3.2 Implications for indicator computation

Now that we got the stochastic viability kernel $\text{Viab}(\beta, \tau)$ for the optimal strategy of 33, dynamic programming is used a second time to get the probability of resilience. This time it is carried out over a set that includes all the states trajectories can visit before getting back to $\text{Viab}(\beta, \tau)$, and a satisfactory set is $(L^*, P) = [0.1, 1] \times [0, 4]$. The proof from A ensures that we find the strategy that maximizes the probability of resilience, and yet again, the solution described through equation (33) is found to apply for all the considered states, and whatever the date and time horizon. Nevertheless, through this section, recall that only $T = 30$ will be used.

If the lake is a hysteretic one, its state is eventually going to become ε-robust after L^* has been brought down to $L^*_{\min}$. This is suggested by Figure 10, which showcases two possibilities for a system that has been deprived of its ε-robustness by a perturbation. It can enter the stochastic viability kernel after a few time steps, or be attracted towards the eutrophic regime. In the latter case, the lake will remain eutrophic for several time steps, so that ε-robustness can only be restored after a much longer time. Thus, the cumulative probability of reaching $\mathrm{Viab}(0.99, 10^3)$ shifts from 0 to 1 in two distinct phases, the first one taking place in the first few time steps and five to six time steps before the second starts (Figure 11). Besides, this figure showcases how the probability of reaching that set in a few time steps dramatically decreases when the post-perturbation state gets farther from its boundary.

Yet for all the considered states, ε-robustness is achieved with unit probability by the horizon $T = 30$, so that the system is resilient when this horizon is considered. One can then choose to use indicators such as those defined in Section 3.4. They can be for instance the mean duration as computed in (20) (Figure 12), or if a worst-case approach is taken, the maximal time needed to become ε-robust with 99% confidence like in equation (21) (Figure 13). Unsurprisingly, both indicators quickly increase as the starting state gets further from the boundary of the stochastic viability kernel, but the increase of the mean time $E(t|f^*, x)$ is less abrupt than that of the maximal $t_\alpha(x)$. This shows that the maximal time needed to get back to a ε-robust set can top 20 time steps when the trajectory involves leaving the set of desirable states and getting into the eutrophic regime. Thus, integrating potential management actions and uncertainty into the model confirms the fact that restoration of the lake's robustness can be a long and costly process in the hysteretic case.

From the maps of Figures 12 and 13, which give the values of indicators for post-perturbation states, one can deduce the impact of different perturbations for a given initial state. In this case study, the perturbation is applied to the system as a one-time jump in the value of one of its state variables. The impact of perturbation amplitude on different indicators can then be observed through curves such as Figure 14, drawn for a system at $(0.5, 0.5)$ before the perturbation. The boundary of the stochastic viability kernel is a threshold beyond which the perturbation can affect the recovery time of the system, which illustrates the interest of keeping its state as far inside this boundary as possible.

Resilience for the irreversible case ($R = 2$) is worth investigating, all other variables keeping the same values. By definition, the system may get stuck in a eutrophic state with no possibility of switching back to the oligotrophic state. Recovery, and therefore resilience, can not guaranteed at any T, which hampers the relevance of using the same indicators as in the hysteretic case. Yet, the probability of resilience $\mathbb{P}_{\max}$ keeps being relevant, and can be computed for $T = 30$ (Figure 15). The feedback map keeps being f^* at each time step, and the horizon is long enough for the system to either recover or switch for good towards an undesirable regime. Like in Figure 14, one can evaluate the impact of a given perturbation at any given pre-perturbation state, using this time the probability of resilience (Figure 16). This illustrates how the proposed framework can foster the use of different yet relevant indicators in distinct situations.

5 Discussion

In the stochastic viability framework, the policy is only optimal with respect to the horizon T. When the feedbacks f^* are computed for a given horizon T, the probabilities of resilience computed with f^* for $t < T$ are not necessarily the maximum probability of reaching $\text{Viab}(\beta, \tau)$ after t time steps. The policy is also only optimal for the objective of maximizing the probability of resilience. In the work by [26], the most resilient trajectory is the one that besides reaching the viability kernel in due time, also minimizes a cost function. In the stochastic case, there are many possible trajectories, some of which may not reach the viability kernel. Thus, minimizing the cost on the trajectory may sometimes conflict with maximizing the probability to reach the viability kernel, and one should define a trade-off between these objectives. Viability theory allows for the exploration of trade-offs in the latter case under the so-called framework of co-viability [14], but without addressing the issue of resilience. Exploring the integration of resilience computations in a co-viability framework could be a natural and useful continuation of this work.

However, in the case of lake eutrophication, the controls are independent on the horizon, so that the feedback map is the same at every date. This means that changing the management horizon will not affect the immediate decisions to be taken. Besides, the map is the same whether ones aims at optimizing the resilience of the system or its ε-robustness. Such a simplification is due to the fact that one always has interest in reducing L^*. Many other systems studied through viability theory in the deterministic case have the same characteristic: controls allow for changing a state variable which increase, or decrease, always leads to enhancing the resilience of the system (see examples throughout [11]). Ecological systems where the same feedback map applies whatever the horizon might in fact be common.

Yet, one should keep in mind that the model exposed in this paper is only a theoretical toy. It is a well-known benchmark model of a nonlinear ecological system, and can be used to explore some aspects of such systems. One of its advantages is that of being two-dimensional, which makes its dynamic easier to understand (Figures 5 and 10) and thus to showcase concepts and methodological developments. However, this model is not meant to be used for assessing resilience of actual lakes, for a number of reasons, some of which can be briefly described here. First, it accounts for mud dynamics only implicitly, through the recycling term $P^q/(1+P^q)$ of equation (24). This term governs the transition from an oligotrophic to a eutrophic state, so that mud dynamics should be explicitly accounted for in any model aimed at being used for decision-making. Second, available policies generally have an unknown impact, so that uncertainty on the control should be included in the model, and they are also not constant in time. For instance, one can expect a policy response to not have immediate effects, and such delays are not taken into account in our hypothetical model. Third, the economic interests of farmers are only reflected by the existence of a bound $L^*_{\min} > 0$ to K, and the optimal policy involves reducing L^* until it reaches that value. Under such a formulation, $L^*_{\min}$ becomes the quantity of phosphorus farmers usually discharge into the lake, even though it may not be acceptable for them to maintain their use at that level for a prolonged period of time.

Last but not least, dynamic programming algorithms work in theory for any system described by equation (14), which would make this method very general. Its area of application, however, is limited in practice by the curse of dimensionality, which causes both the required computational time and

memory to increase exponentially with the dimension of the state space. Yet, the notion of stochastic viability kernel, and that of resilience of a system based on a policy, are not defined through dynamic programming, which is merely a good way to compute them. It could be very interesting to explore in future work how they can be generalized to systems of higher dimensionality, even if such developments may render the computations more approximate.

6 Conclusions

This work presented a framework for defining and computing resilience of stochastic controlled dynamical systems, using stochastic viability kernels and dynamic programming. This allows for the incorporation of uncertainty so that computations can explicitly account for the capacity of a system to cope with adverse events. The two main advantages of the viability framework of resilience prove all their utility in the uncertain case. On the one hand, attractors cease to be fixed points when the dynamics is stochastic, yet viability tools such as stochastic viability kernels can still be defined. On the other hand, the possibility to use this framework to dynamically optimize the management policies while computing the resilience remains in the stochastic case.

The presented framework does not claim to give a single, standard formula for computing resilience, because what can be computed in practice is only a resilience indicator. Rather, it proposes a set of possibilities for deriving general indicators that can fit different applicative contexts. The case of the lake illustrates how this framework allows for the computation of different indicators for resilience depending on the value of the lake parameter R. In the end, the indicators given in this work are generic examples: they may not be the most relevant one in a given case, but they can be built upon to define such a relevant indicator.

Acknowledgments

This work has been supported by grants from Irstea and Région Auvergne.

A Relationship between V_d, V_s and resilience.

In the work by [13], given a time horizon T, one is interested in keeping a system described by uncertain dynamics such as g of equation (14) in a constraint set $A(t)$ for all dates $t < T$. For simplicity, we also suppose that $A(t)$ represents a discretezed set of points. A dynamic programming algorithm is proposed and uses the following value function:

$$\begin{cases} V(T,x) = \begin{cases} 1 & \text{if} \quad x \in A(T) \\ 0 & \text{if} \quad x \notin A(T) \end{cases} \\ \forall t \in [0, T-1],\ V(t,x) = \max\limits_{u_t \in U(x)} \left(\sum\limits_{y \in A(t)} \mathbb{P}(g_\varepsilon(x, u_t) = y)\ V(t+1, y) \right) \end{cases} \tag{34}$$

The main result of their work is that this value function allows for finding the feedbacks that maximizes the odds of the system to remain in $A(t)$ at all dates $t < T$. In particular, setting $T = \tau$ and $A(t) = K$ at all dates yields the value function used to compute stochastic viability kernels. $x \in \text{Viab}(\beta, \tau)$ is then equivalent to $V(0, x) > \beta$.

Equation (34) can equivalently be written by considering the horizon $T - t$ instead of the date t:

$$\begin{cases} V(0,x) = \begin{cases} 1 & \text{if} \quad x \in A(0) \\ 0 & \text{if} \quad x \notin A(0) \end{cases} \\ \forall t \in [1, T],\ V(t,x) = \max\limits_{f_t(x) \in U(x, T-t)} \left(\sum\limits_{y \in A(t)} \mathbb{P}(g_\varepsilon(x, f_t(x)) = y)\ V(t-1, y) \right) \end{cases} \tag{35}$$

Setting $A(0)$ as the viability kernel and $A(t)$ as the entire state space for any horizon $t > 0$, and placing ourselves in the case where no uncertainty is present, V becomes the value function V_d. Then V_d is the optimal probability of the problem of reaching the viability kernel by T, and the set of states x such that $V_d(T, x) = 1$ is the set of states which are solution to this problem, which corresponds to the definition of the resilience basin $\text{Res}(T)$. The same applies for any $t \leq T$.

Likewise for V_s, one can set $A(0) = \text{Viab}(\beta, \tau)$ and $A(t)$ as the entire state space for $t < T$ in equation (35). Furthermore, to express that we are not interested in this computation by what might happen after the recovery is complete and the system reaches the viability kernel, we need to apply the results from [13] to modified dynamic in which all the points within $\text{Viab}(\beta, \tau)$ are fixed whatever the control and uncertainty ($\forall u, \varepsilon, g(x, u, \varepsilon) = x$). Then, V_s is the optimal probability of the problem of being in $\text{Viab}(\beta, \tau)$ at date T under the modified dynamics, or to reach it before T under the original dynamics. Hence, for any state x, $\mathbb{P}_{\text{Res}}(T, x) = V_s(T, x)$ because both entail optimizing the same probability over the same time frame.

B Dimensional analysis of the lake problem

Recall the continuous time equation (23):

$$\frac{dp}{d\theta} = -b\,p + l + r\frac{p^8}{p^8 + m^8} \tag{36}$$

We want to reduce the number of parameters by introducing the following dimensionless variables and parameters:

$$P = \frac{p}{m} \qquad L = \frac{l}{b\,m} \qquad R = \frac{r}{b\,m} \qquad dt = b\,d\theta \tag{37}$$

so that dividing equation (36) by $(b\,m)$ yields equation (24), effectively reducing the number of parameters from three to one.

C Derivation of the discrete time equation (27)

We start from the continuous time equation (24), integrate uncertainty and controls before deriving the discrete time equivalent. The amount of time between dates t and $t+1$ is noted Δt, and is supposed constant between any pair of consecutive dates.

C.1 Uncertainty

The continuous time equivalent for equation (25) is:

$$L(t) = L^*(t) + \omega(t) \tag{38}$$

where ω is a brownian motion such that its pdf after a time interval Δt is $\mathcal{N}(0, \sigma)$.

C.2 Controls

In the continuous time formulation of [26], the introduction of a limited capacity of action in equation (26) is mathematically translated into the possibility to choose the temporal derivative of L^* within a bounded range, noted U_c in continuous time. This can be written as:

$$\begin{cases} \frac{dL^*}{dt} = L^* + u \\ u \in U_c(x) = \left[\frac{U_c^{\min}(x)}{\Delta t}, \frac{U_c^{\max}(x)}{\Delta t}\right] \end{cases} \tag{39}$$

C.3 Derivation of the discrete-time evolution equation

Integrating uncertainty, equation (24) becomes:

$$\frac{dP}{dt} = -P + L^* + \omega + R\frac{P^8}{p^8 + m^8} \tag{40}$$

Let us introduce δt such that $\Delta t = j\delta t$ with $j \in \mathbb{N}$, and use the smaller δt for the Euler approximation of equation (24) so as to minimize the computational error. For any $k \in [0, j-1]$ and a given value of u, we get:

$$\begin{cases} P_{t+(k+1)\delta t} = P_{t+k\delta t} + \delta t\left(-P_{t+k\delta t} + L^*_{t+k\delta t} + \frac{\sigma}{\Delta t}\varepsilon_t + R\frac{P^8_{t+k\delta t}}{P^8_{t+k\delta t} + 1}\right) \\ L^*_{t+(k+1)\delta t} = L^*_{t+k\delta t} + u\frac{\delta t}{\Delta t} \end{cases} \tag{41}$$

and the latter equation can be iterated j times over to get equation 27:

$$P_{t+1} = g(P_t, L_t^*, u_t, \varepsilon_t) \tag{42}$$

where the parameters R and σ are made implicit. Throughout this paper we used $j = 10$.

References

[1] W. N. Adger. Social and ecological resilience: are they related? *Progress in Human Geography*, 24(3):347–364, 2000.

[2] W. N. Adger, T. P. Hugues, C. Folke, S. R. Carpenter, and J. Rockström. Social-ecological resilience to coastal disasters. *Science*, 309:1036–1039, 2005.

[3] C. R. Allen, J. J. Fontaine, K. L. Pope, and A. S. Garmestani. Adaptive management for a turbulent future. *Journal of Environmental Management*, 92:1339–1345, 2011.

[4] J. M. Anderies, M. A. Janssen, and B. H. Walker. Grazing management, resilience, and the dynamics of a fire-driven rangeland system. *Ecosystems*, 5:23–44, 2002.

[5] F. S. Brand and K. Jax. Focusing on the meaning(s) of resilience: Resilience as a descriptive concept and a boundary object. *Ecology and Society*, 12(1), 2007.

[6] S. Carpenter, B. Walker, J. M. Anderies, and N. Abel. From metaphor to measurement: Resilience of what to what? *Ecosystems*, 4:765–781, 2001.

[7] S. R. Carpenter, D. Ludwig, and W. A. Brock. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications*, 9(3):751–771, 1999.

[8] M. De Lara and L. Doyen. *Sustainable Management of Natural Resources*. Springer, 2008.

[9] M. De Lara and V. Martinet. Multi-criteria dynamic decision under uncertainty: A stochastic viability analyis and an application to sustainable fishery management. *Mathematical Biosciences*, 217:118–124, 2009.

[10] G. Deffuant, L. Chapel, and S. Martin. Approximating viability kernel with support vector machines. *IEEE Transactions on Automatic Control*, 52(5):933–937, 2007.

[11] G. Deffuant and N. Gilbert, editors. *Viability and resilience of complex systems*. Kluwer Academic Publishers, London, 2011.

[12] L. Doyen and C. Béné. Sustainability of fisheries through marine reserves: a robust modeling analysis. *Journal of Environmental Management*, 69:1–13, 2003.

[13] L. Doyen and M. De Lara. Stochastic viability and dynamic programming. *Systems and Control Letters*, 59:629–634, 2010.

[14] L. Doyen, O. Thébaud, C. Béné, V. Martinet, S. Gourguet, M. Bertignac, S. Fifas, and F. Blanchard. A stochastic viability approach to ecosystem-based fisheries management. *Ecological Economics*, 75:32–42, 2012.

[15] C. Folke, S. Carpenter, T. Elmqvist, L. Gunderson, C. S. Holling, B. Walker, J. Bengtsson, F. Berkes, J. Colding, K. Danell, M. Falkenmark, L. Gordon, R. Kasperson, N. Kautsky, A. Kinzig, S. Levin, K.-G. Mäler, F.Moberg, L. Ohlsson, P. Olsson, E. Ostrom, W. Reid, J. Rockström, H. Savenije, and U. Svedin. Resilience and sustainable development: Building adaptive capacity

in a world of transformations. Technical report, Scientific Background Paper on Resilience for the process of The World Summit on Sustainable Development on behalf of The Environmental Advisory Council to the Swedish Government, April 2002.

[16] C. S. Folke, S. Carpenter, B. Walker, M. Scheffer, T. Elmqvist, L. Gunderson, and C. S. Holling. Regime shifts, resilience, and biodiversity in ecosystem management. *Annual Review of Ecology, Evolution, and Systematics*, 35:557–581, 2004.

[17] V. Grimm and C. Wissel. Babel, or the ecological stability discussions: an inventory and analysis of terminology and a guide for avoiding confusion. *Oecologia*, 109:323–334, 1997.

[18] L. H. Gunderson. Ecological resilience – in theory and application. *Annual Review of Ecology and Systematics*, 31:425–439, 2000.

[19] V. Guttal and C. Jayaprakash. Impact of noise on bistable ecological systems. *Ecological Modelling*, 201(3-4):420–428, 2007.

[20] T. Hashimoto, J. R. Stedinger, and D. P. Loucks. Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. *Water Resources Research*, 18(1):14–20, 1982.

[21] C. S. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology, Evolution, and Systematics*, 4:1–23, 1973.

[22] K. H. Johnson, K. A. Vogt, H. J. Clark, O. J. Schmitz, and D. J. Vogt. Biodiversity and the productivity and stability of ecosystems. *Trends in Ecology and Evolution*, 11(9):372–377, September 1996.

[23] M. Lesnoff, C. Corniaux, and P. Hiernaux. Sensitivity analysis of the recovery dynamics of a cattle population following drought in the sahel region. *Ecological Modelling*, 232:28–39, 2012.

[24] D. Ludwig, S. Carpenter, and W. Brock. Optimal phosphorus loading for a potentially eutrophic lake. *Ecological Applications*, 13(4):1135–1152, 2003.

[25] C. W. MacGillivray and J. P. Grime. Testing predictions of the resistance and resilience of vegetation subjected to extreme events. *Functional Ecology*, 9:640–649, 1995.

[26] S. Martin. The cost of restoration as a way of defining resilience: a viability approach applied to a model of lake eutrophication. *Ecology and Society*, 9(2), 2004.

[27] J.-C. Péreau, L. Doyen, L. R. Little, and O. Thébaud. The triple bottom line: Meeting ecological, economic and social goals with individual transferable quotas. *Journal of Environmental Economics and Management*, 63:419–434, 2012.

[28] P. Saint-Pierre. Approximation of the viability kernel. *Applied Mathematics and Optimization*, 29:187–209, 1994.

[29] B. Walker, S. Carpenter, J. Anderies, N. Abel, G. Cumming, M. Janssen, L. Lebel, J. Norberg, G. D. Peterson, and R. Pritchard. Resilience management in social-ecological systems: a working hypothesis for a participatory approach. *Conservation Ecology*, 6(1):14, 2002.

[30] B. Walker, C. S. Holling, S. R. Carpenter, and A. Kinzig. Resilience, adaptability and transformability in social-ecological systems. *Ecology and Society*, 9(2):5, 2004.

[31] B. K. Williams. Adaptive management of natural resources – framework and issues. *Journal of Environmental Management*, 92:1346–1353, 2011.

List of Figures

1 Illustration of the notions of strategy and time horizon (equations (2) through (5)). . . 27

2 Constraint set, viability kernel and some resilience basins in the original viability framework of [26] and [11]. There is one possible trajectory for a given control strategy: post-perturbation states A and B are resilient while C is not. 27

3 Constraint set, stochastic viability kernel and level sets of the probability of resilience given a horizon T. Some possible trajectories starting from the post-perturbation state A are represented: (a) and (b) are resilient trajectories while (c) and (d) are not because they are outside $\mathrm{Viab}(\beta, \tau)$ after T time steps. 28

4 Stable (continuous line) and unstable (dashed line) attractors for the lake eutrophication problem. 29

5 Two trajectories simulated in the (L, P) plane using $\sigma = 0.1$ and with a constant value of $L^* = 0.63$. Both start at $x_0 = (0.6, 0.63)$, which is an oligotrophic stable equilibrium. The dates are $t = 0, \ldots, 15$, and the final date is signaled with a cross. 30

6 The 95% confidence interval (CI) for the position of the attractors is represented in grey in the case $\sigma = 0.1$, as well as the 95% CI for the state x_5 after 5 times steps, if the initial state is $x_0 = (0.6, 0.63)$. 31

7 Stochastic viability kernels for $\beta = 0.99$ and different values of σ and τ, reminding that for $\tau_1 < \tau_2$, $\mathrm{Viab}(\beta, \tau_2)$ is a subset of $\mathrm{Viab}(\beta, \tau_1)$. For reference the viability kernel ($\sigma = 0$) is on the left side of the grey line. 32

8 Probability of resilience at $T = 30$ and entry time t_α for strategy $f^{(1)}$. By default, $t_\alpha(x) = 30$ when the probability of resilience is smaller than 0.99 over the management horizon. 33

9 Probability of resilience at $T = 30$ and entry time t_α for strategy $f^{(2)}$. By default, $t_\alpha(x) = 30$ when the probability of resilience is smaller than 0.99 over the management horizon. 34

10 Six trajectories towards the ε-robust set in the (L^*, P) plane, simulated with $\sigma = 0.1$, $M = 0.05$ and a post-perturbation state $(0.8, 0.5)$. 35

11 Cumulative probability of reaching an ε-stable state in t time steps, for different post-perturbation values of the mean phosphorus input L^*, and a phosphorus concentration $P = 0.5$. 35

12 The resilience-related indicator $E(t|f^*, x)$ for a hysteretic lake ($R = 1.2$) with $M = 0.05$, $\sigma = 0.1$ and $T = 30$. 36

13 The resilience-related indicator $t_\alpha(x)$ ($\alpha = 0.99$) for a hysteretic lake ($R = 1.2$) with $M = 0.05$, $\sigma = 0.1$ and $T = 30$. 37

14 Impact of the amplitude of a perturbation on resilience-related indicators for a hysteretic lake, if the system is at $(0.5, 0.5)$ before the perturbation. 38

15 $\mathbb{P}_{\mathrm{Res}}(T, x)$ for an irreversible lake ($R = 2$), and $T = 30$. 39

16 For an irreversible lake which state is $(0.5, 0.5)$ before the perturbation, possible impacts of a perturbation on the probability of resilience. 40

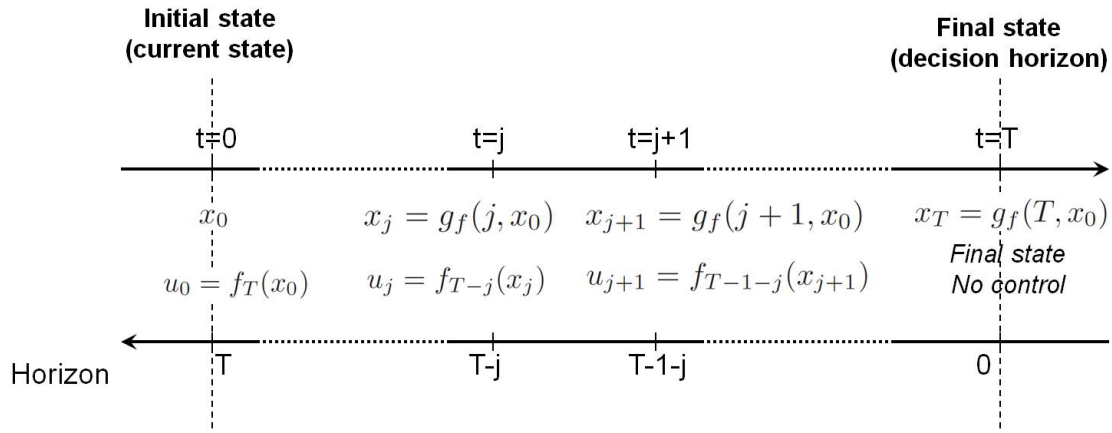

Figure 1: Illustration of the notions of strategy and time horizon (equations (2) through (5)).

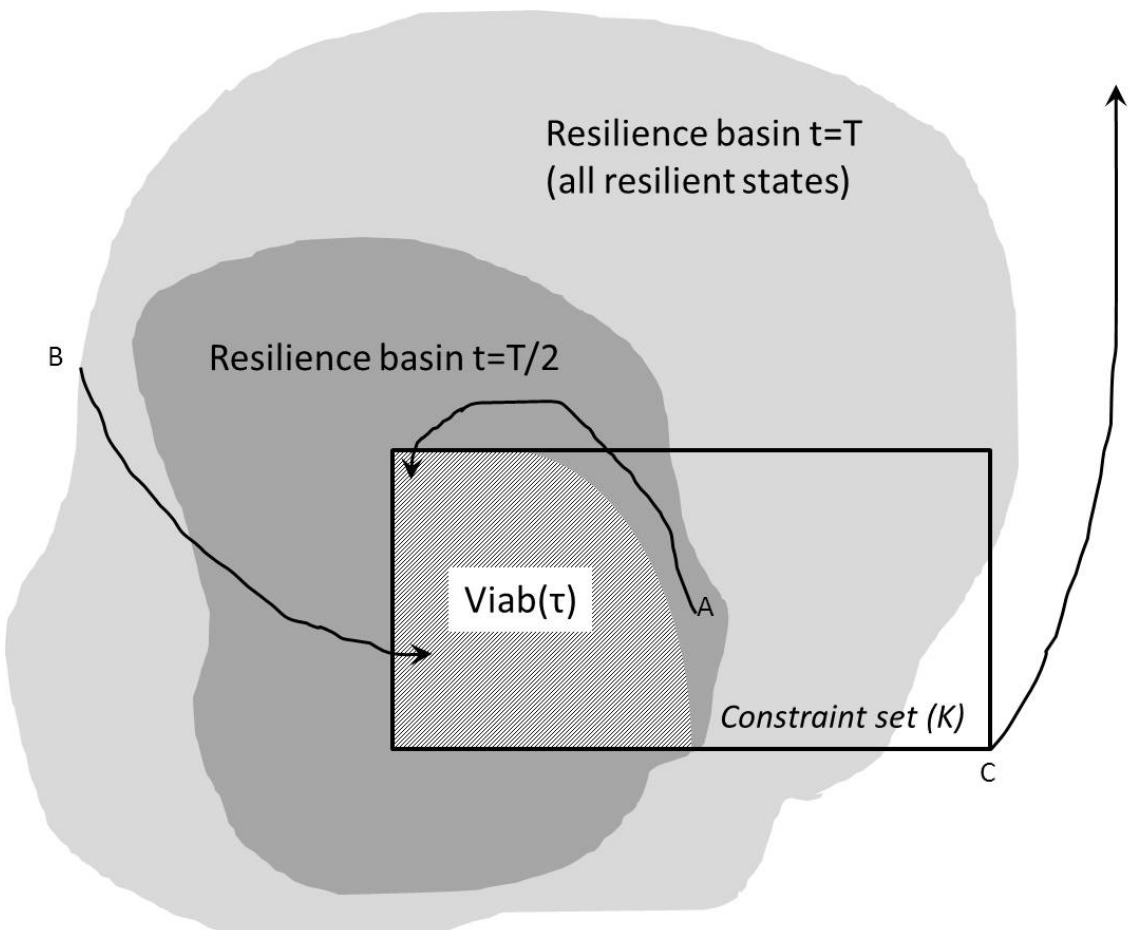

Figure 2: Constraint set, viability kernel and some resilience basins in the original viability framework of [26] and [11]. There is one possible trajectory for a given control strategy: post-perturbation states A and B are resilient while C is not.

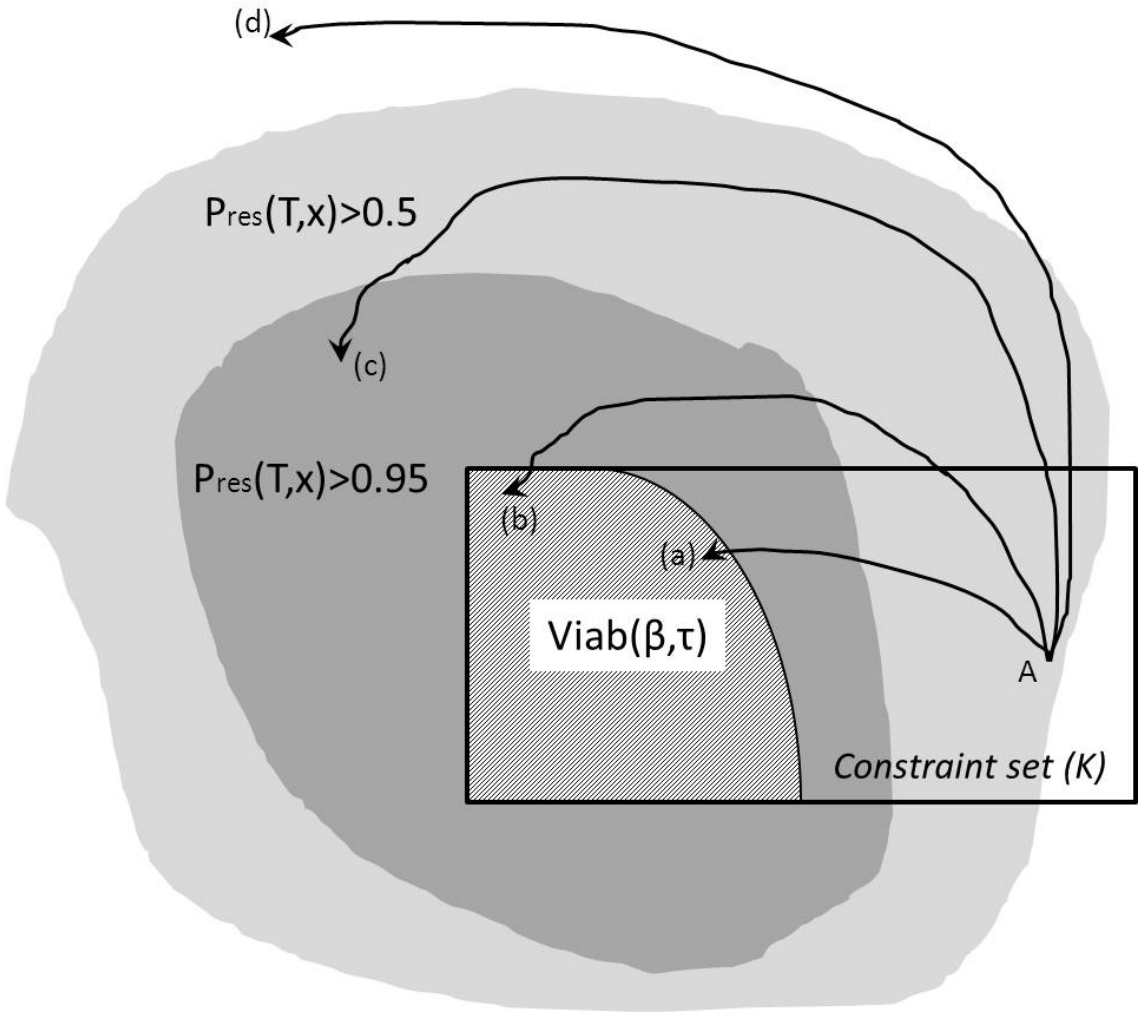

Figure 3: Constraint set, stochastic viability kernel and level sets of the probability of resilience given a horizon T. Some possible trajectories starting from the post-perturbation state A are represented: (a) and (b) are resilient trajectories while (c) and (d) are not because they are outside $\mathrm{Viab}(\beta, \tau)$ after T time steps.

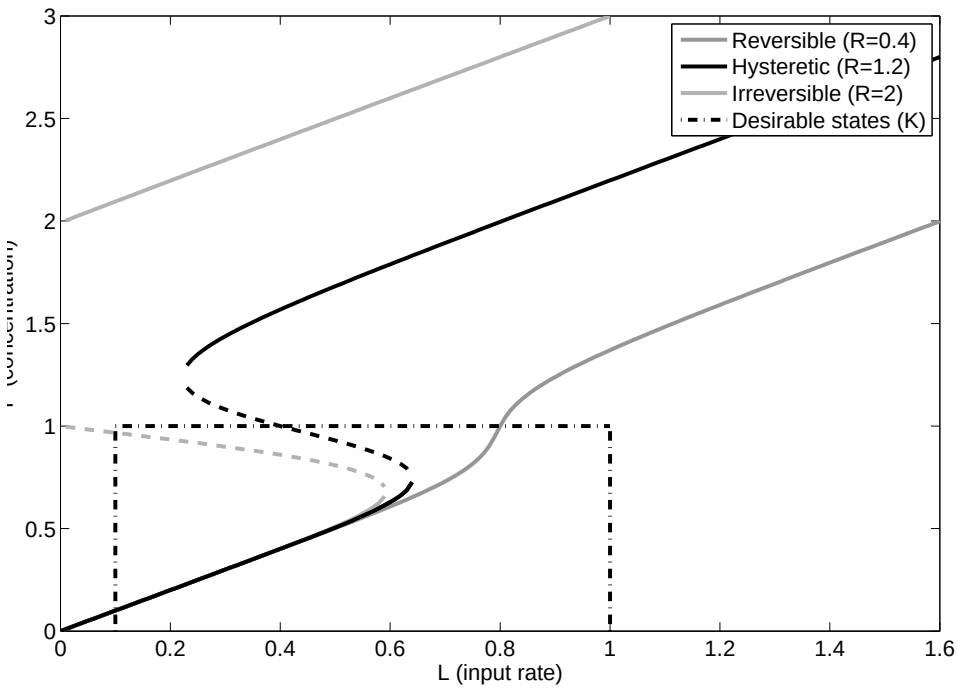

Figure 4: Stable (continuous line) and unstable (dashed line) attractors for the lake eutrophication problem.

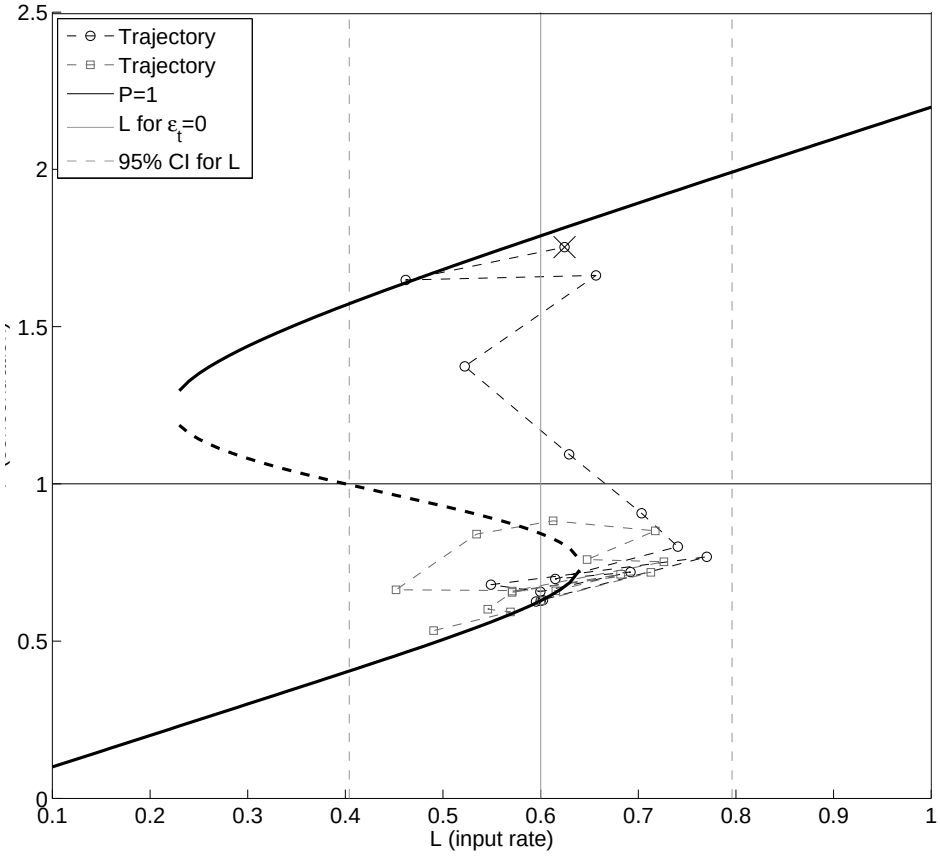

Figure 5: Two trajectories simulated in the (L, P) plane using $\sigma = 0.1$ and with a constant value of $L^* = 0.63$. Both start at $x_0 = (0.6, 0.63)$, which is an oligotrophic stable equilibrium. The dates are $t = 0, \ldots, 15$, and the final date is signaled with a cross.

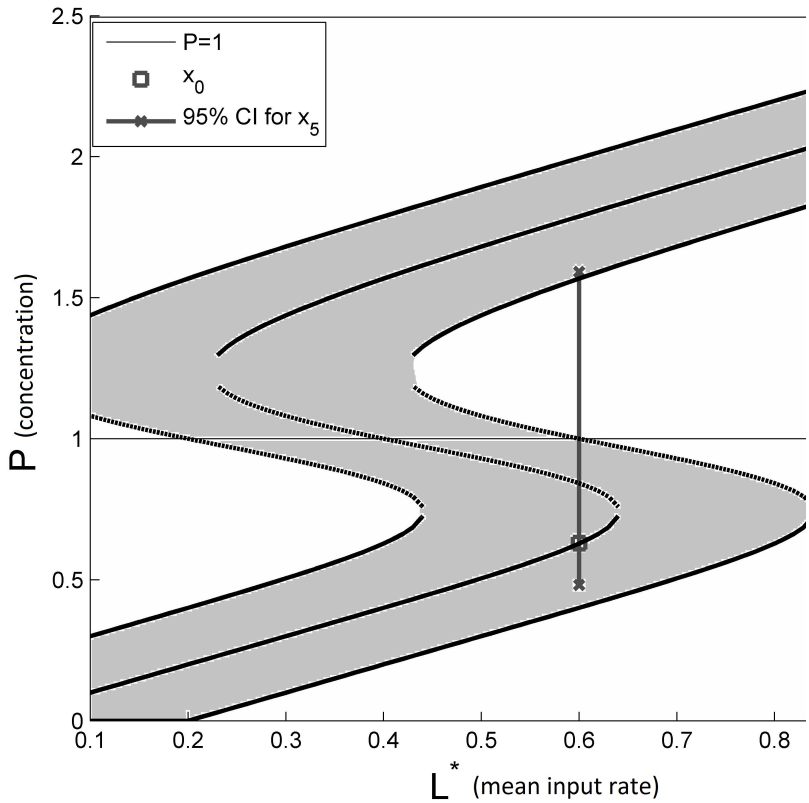

Figure 6: The 95% confidence interval (CI) for the position of the attractors is represented in grey in the case $\sigma = 0.1$, as well as the 95% CI for the state x_5 after 5 times steps, if the initial state is $x_0 = (0.6, 0.63)$.

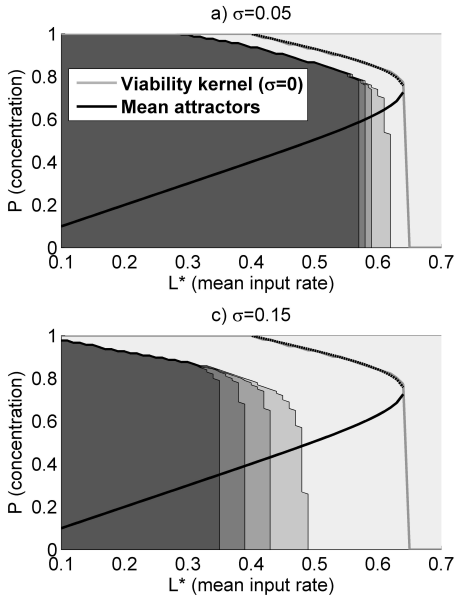

Figure 7: Stochastic viability kernels for $\beta = 0.99$ and different values of σ and τ, reminding that for $\tau_1 < \tau_2$, $\mathrm{Viab}(\beta, \tau_2)$ is a subset of $\mathrm{Viab}(\beta, \tau_1)$. For reference the viability kernel ($\sigma = 0$) is on the left side of the grey line.

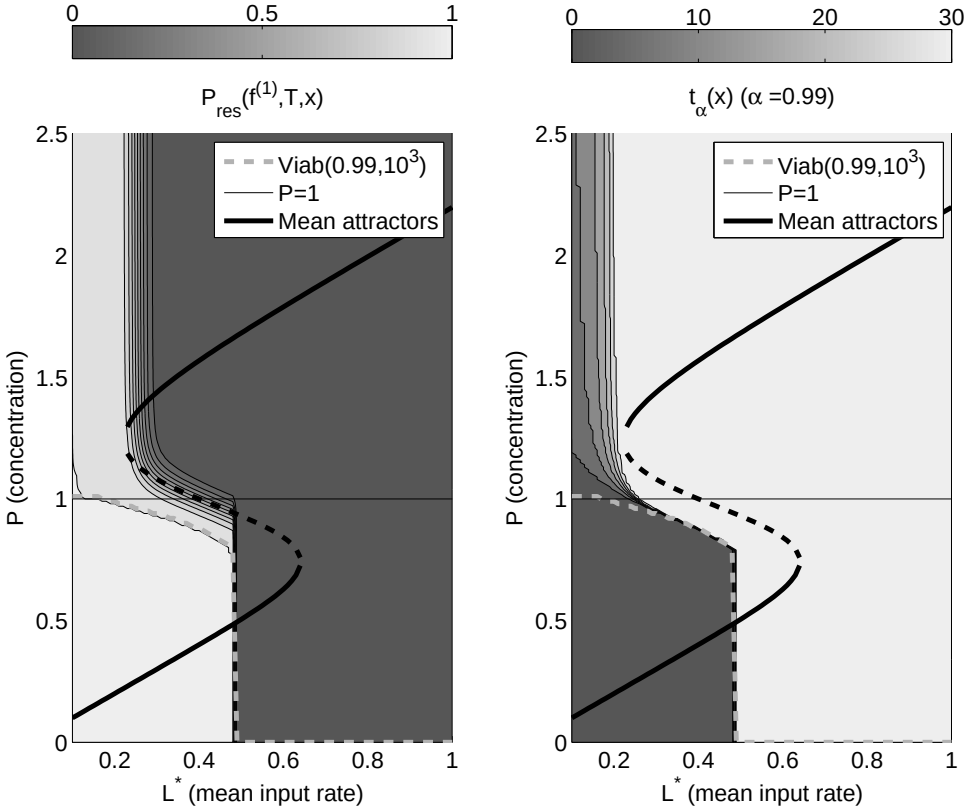

Figure 8: Probability of resilience at $T = 30$ and entry time t_α for strategy $f^{(1)}$. By default, $t_\alpha(x) = 30$ when the probability of resilience is smaller than 0.99 over the management horizon.

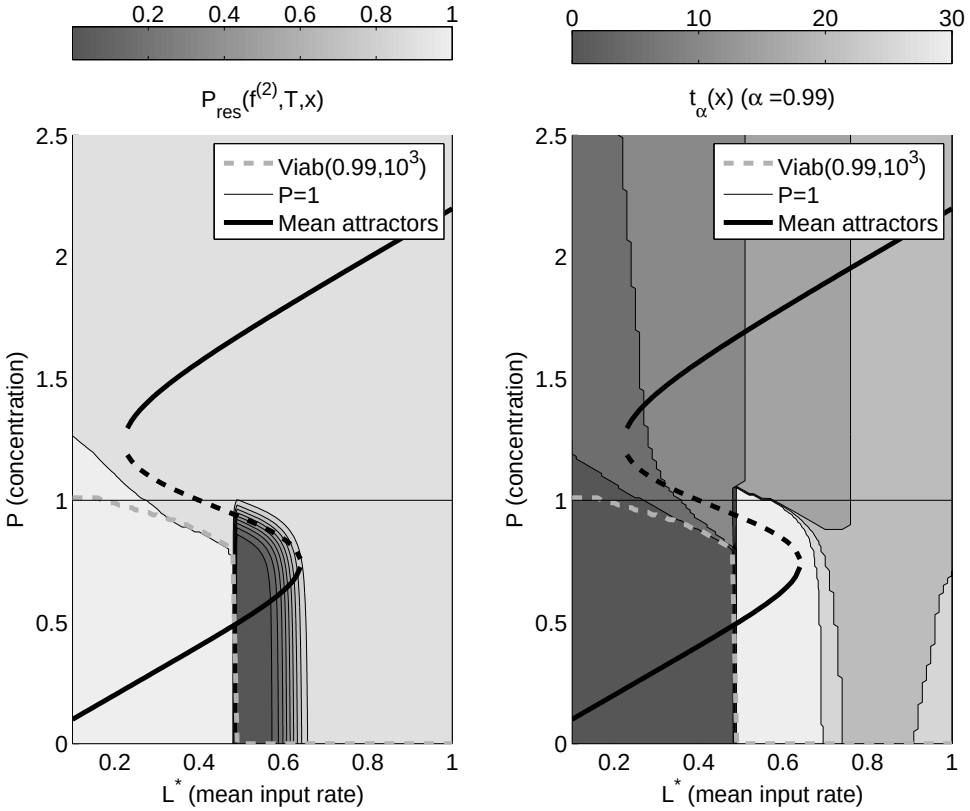

Figure 9: Probability of resilience at $T = 30$ and entry time t_α for strategy $f^{(2)}$. By default, $t_\alpha(x) = 30$ when the probability of resilience is smaller than 0.99 over the management horizon.

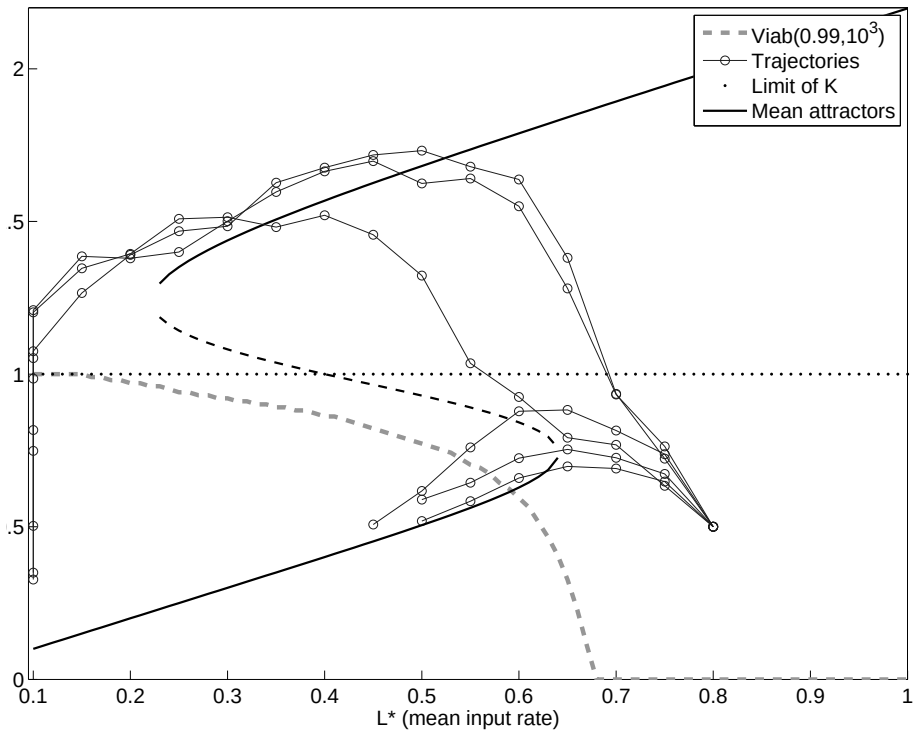

Figure 10: Six trajectories towards the ε-robust set in the (L^*, P) plane, simulated with $\sigma = 0.1$, $M = 0.05$ and a post-perturbation state $(0.8, 0.5)$.

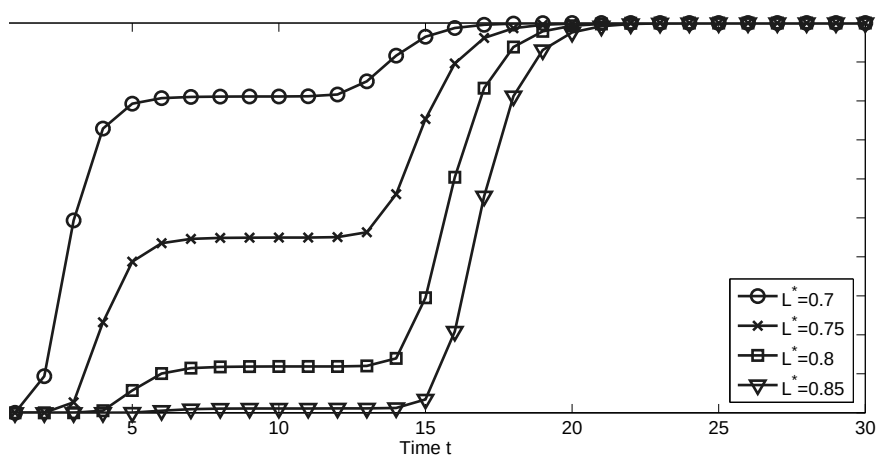

Figure 11: Cumulative probability of reaching an ε-stable state in t time steps, for different post-perturbation values of the mean phosphorus input L^*, and a phosphorus concentration $P = 0.5$.

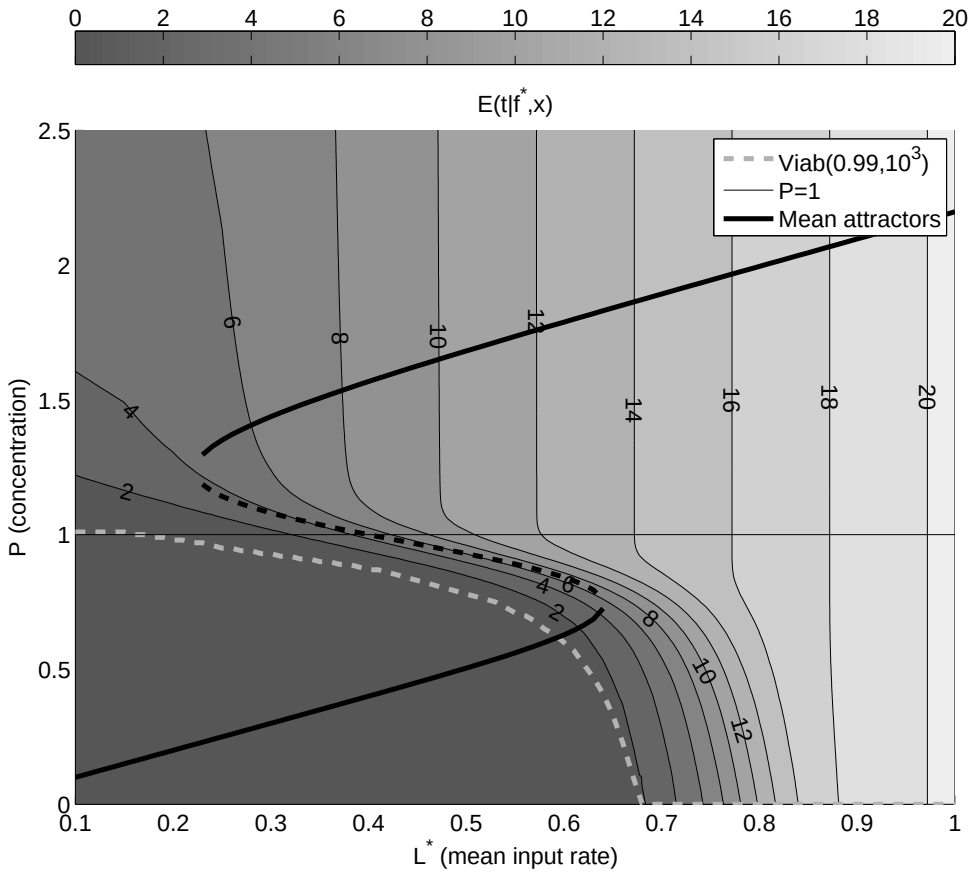

Figure 12: The resilience-related indicator $E(t|f^*, x)$ for a hysteretic lake ($R = 1.2$) with $M = 0.05$, $\sigma = 0.1$ and $T = 30$.

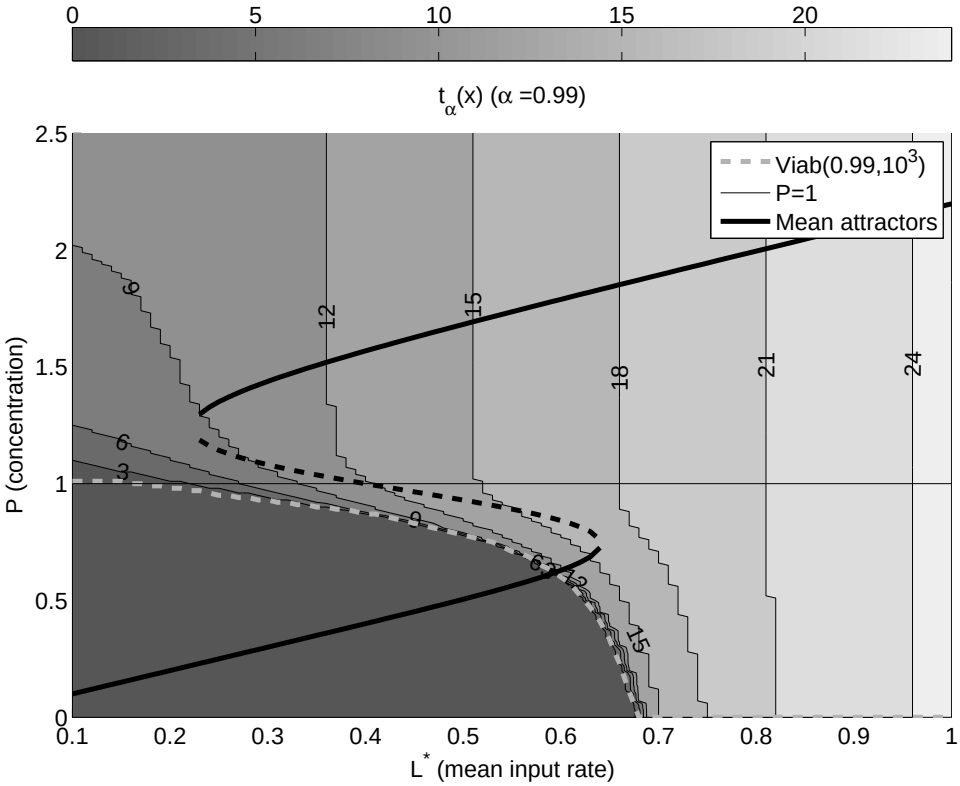

Figure 13: The resilience-related indicator $t_\alpha(x)$ ($\alpha = 0.99$) for a hysteretic lake ($R = 1.2$) with $M = 0.05$, $\sigma = 0.1$ and $T = 30$.

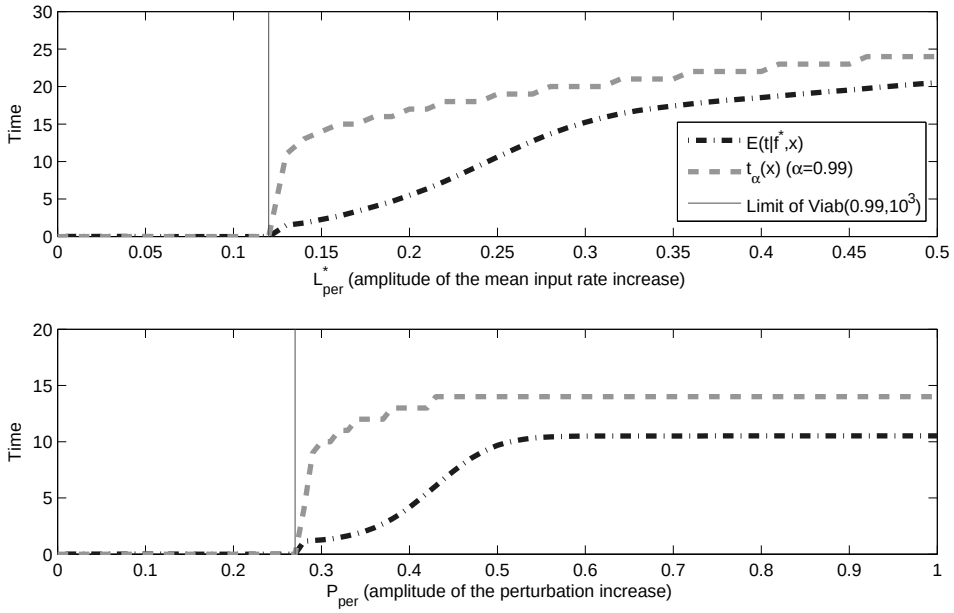

Figure 14: Impact of the amplitude of a perturbation on resilience-related indicators for a hysteretic lake, if the system is at $(0.5, 0.5)$ before the perturbation.

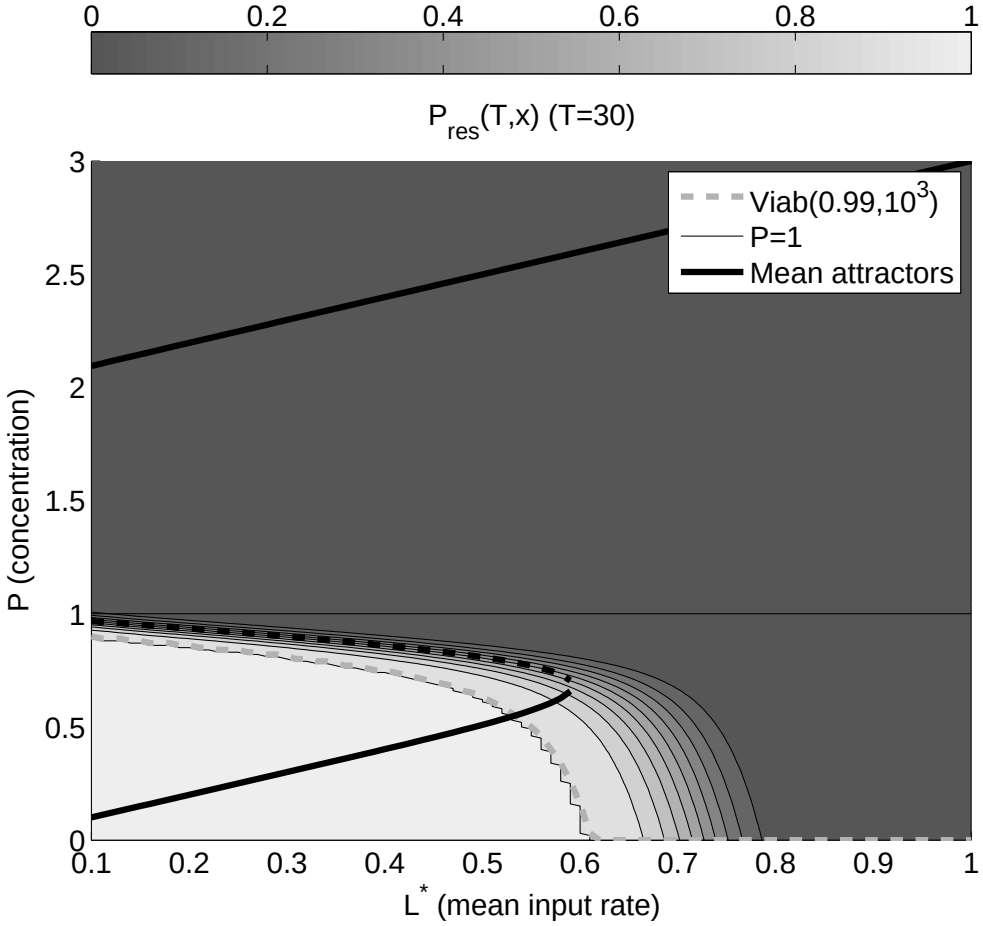

Figure 15: $\mathbb{P}_{\mathrm{Res}}(T, x)$ for an irreversible lake ($R = 2$), and $T = 30$.

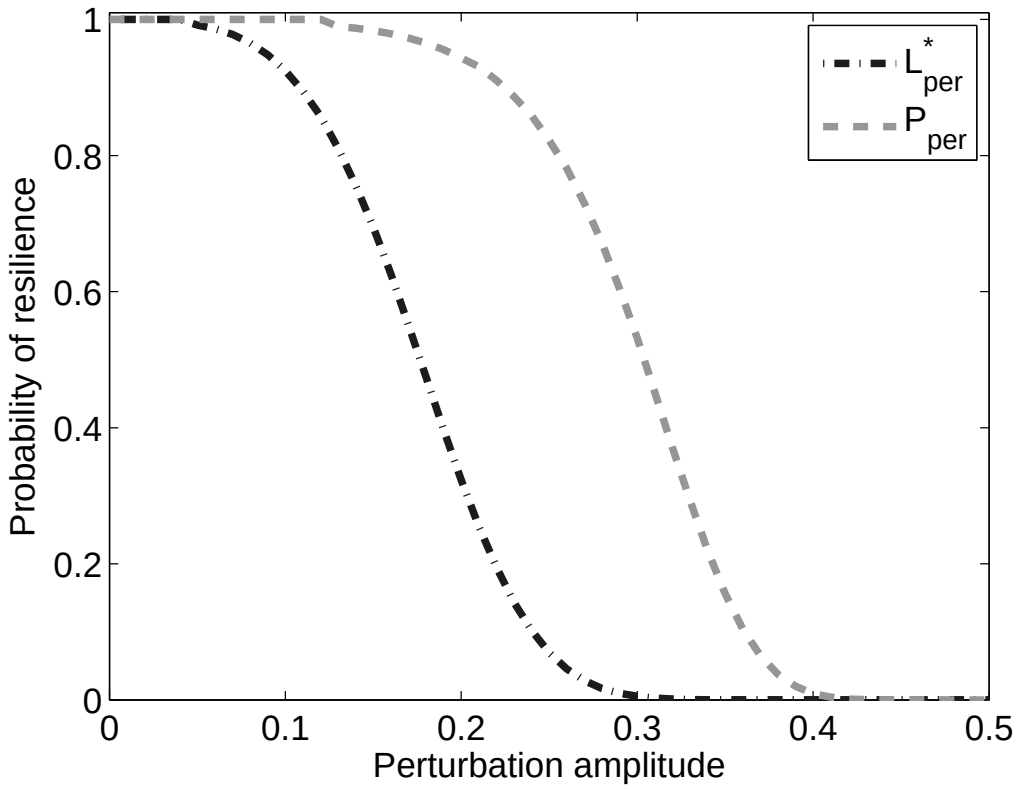

Figure 16: For an irreversible lake which state is $(0.5, 0.5)$ before the perturbation, possible impacts of a perturbation on the probability of resilience.