Introduction to Theoretical Ecology

Natal, 2011

Objectives

After this week:

- The student understands the concept of a biological system in equilibrium and knows that equilibria can be stable or unstable.
- The student understands the basics of how coupled differential equations can be analyzed graphically, including phase plane analysis and nullclines.
- The student can analyze the stability of the equilibria of a one-dimensional differential equation model graphically.
- The student has a basic understanding of what a bifurcation point is.
- The student can relate alternative stable states to a 1D bifurcation plot (e.g. catastrophe fold).

Study material / for further study:

This text

Scheffer, M. 2009. Critical Transitions in Nature and Society, Princeton University Press, Princeton and Oxford.

Scheffer, M. 1998. Ecology of Shallow Lakes. 1 edition. Chapman and Hall, London.

Edelstein-Keshet, L. 1988. Mathematical models in biology. 1 edition. McGraw-Hill, Inc., New York.

Tentative programme (maybe too tight for the exercises)

Monday
9:00-10:30 Introduction Modelling + introduction Forrester diagram + 1D models (stability graphs)
10:30-13:00 GRIND Practical CO_2 chamber - Ethiopian Wolf

Tuesday
9:00-10:00 Introduction bifurcation (Allee effect) and Phase plane analysis (Lotka-Volterra competition)
10:00-13:00 GRIND Practical Lotka-Volterra competition + Sahara

Wednesday
9:00-13:00 GRIND Practical – Sahara (continued) and Algae-zooplankton

Thursday
9:00-13:00 GRIND practical – Algae zooplankton spatial heterogeneity

Friday
9:00-12:00 GRIND practical- Algae zooplankton fish
12:00-13:00 Practical summary/explanation of results - Wrap up

An introduction to models

What is a model?

The word 'model' is used widely in every-day language. It is therefore not surprising that confusion commonly arises around the term. A generic definition of a model is 'a simplified representation of reality'. This includes descriptions in words (the socialistic world model), physical representations (maps, model airplanes) and mathematical models, which we will address in this course. In fact the mathematical models have several things in common with word-models or physical representations. Perhaps the most crucial common aspect is that all such models leave out many details in order to allow us to focus better on some essential aspects. The question which aspects are essential depends on what you want to use the model for. This may differ widely, and consequently many very different models of the same thing (e.g. a lake) may exist, none of which is 'the best' in a general sense. Consider the analogue of the map. Different maps exist for hikers, cars, railways, airplanes, geologists etc. They all represent subsets of the real world selected for different purposes. Maps are useful because they leave out most of the details. They would be unreadable otherwise. In fact, by definition there is only one model, which includes everything: the real world itself.

This immediately identifies the core of what modeling is about: simplifying. The challenge is to do that in the best way, given the goal one has in mind. Broadly speaking, there are two categories of goals that one may aim at with the use of mathematical models: *prediction* and *understanding*. It is widely recognized that these goals require quite different modeling approaches. For prediction of the effects of specific human impacts on environmental quality or animal and plant communities an array of mathematical model types exist. On one end of the range are the *empirical models* such as classic statistical models and various artificial intelligence approaches such as neural networks and expert systems. On the other end there are elaborate *simulation models* that cover the main processes that govern the systems behavior in a mechanistic way. By contrast models that are used in various scientific disciplines to help understanding how systems work are usually very simple. These models which often appear in scientific front-line literature, are sometimes called '*minimal models*' or 'strategic models' and are meant to reveal possible explanations for observed phenomena.

In the following sections we first discuss the scope of different types of predictive models and also argue that a lack of mechanistic insight in how systems work is typically the bottleneck to good prediction. We then turn to a reflection on the fundamental barriers to obtaining a true mechanistic understanding of how nature works. Finally, we discuss the type of models used to develop scientific insight rather than to predict.

Models for prediction

In practice, the task of the environmental scientists often boils down to predicting the effect of different possible management scenarios. This can be on small scales. For instance, it may be important to know how water quality and biodiversity in a lake would benefit from installation of a water treatment plant, which reduces the concentrations of nutrients, and other polluting components of inflowing water. On the other hand models are used to predict the impact of rising CO_2 levels on the climate of the entire earth and on ocean currents.

Empirical models

Somehow the most straightforward way of predictive modeling is the use of empirical (or statistical) relationships.

For instance, by studying many rather deep unvegetated lakes it has been shown that the concentration of phosphorus in the water of the lakes (P) is related to the concentration of the inflowing water (P_i) and the time the water stays in the lake (the so-called hydraulic retention time τ_r) in a predictable way (Vollenweider, 1977):

$$P = c \frac{P_i}{1 + \sqrt{\tau_r}}$$

This empirical 'Vollenweider equation' describes a generic relationship between input and the equilibrium concentration in the lake. Note that the lake concentration is always lower than the input concentration, due to a net loss of phosphorus with sinking particles to the sediment. The model accounts for this effect. It has been used widely to predict the effect of a reduction in nutrient load on the concentration in the lake water. After a transient period the nutrient concentration in the lake should settle to a new equilibrium value depending on the new input concentration and hydraulic retention time in the way described by Vollenweider's equation. However, there are limitations to the model. For instance, shallow lakes in which clarity of the water improves may develop a vegetation of submerged plants, which strongly affects the way phosphorus is partitioned between sediments and water, making the model in its current form useless. The problem here is one of *extrapolation*. We cannot extrapolate the empirical relationship derived from a certain limited set of lakes to a new situation (vegetated shallow lakes).

More serious than the extrapolation problem, is the problem of erroneous assumptions on causality. In the absence of insight in the dominant causal mechanisms, empirical relationships are a highly unreliable basis for managing systems. Of course, this is a common notion, but since it is one of the main caveats in every day applied science, it deserves some repetition. As a caricatural example to highlight the topic once again, consider the relationship between wind and the swinging of trees. One could well imagine that swinging trees actually cause the wind: If they stop swinging the wind stops also. Such a lack of insight in the functioning of the system becomes a problem if one decides to manage the system by cutting down trees in order to reduce the wind. In practice, the caveats of applying empirical models for management purposes are often more tricky than the simple cause-effect exchange in this tree example. The computed values of parameters in statistical models, for instance, are dependent on the correlative structure between the used explanatory variables as well as the non-investigated but also causally related variables (Montgomery and Peck, 1982). This implies that if this correlative structure is affected by a measurement, the model cannot be used for prediction anymore.

The problems of extrapolation and causality apply just as strongly to various sophisticated artificial intelligence approaches that relate 'inputs to outputs'. An example is artificial neural networks (ANN). These are in fact simple simulation models of real networks of neural cells as they appear for instance in the human brain. An ANN adjusts its parameters automatically through a certain algorithm during a 'training process' in which it is offered inputs (e.g. the handwritten image of a letter) and correct outputs (e.g. which letter it is). When offered a new input, a well-trained network gives the correct

output. In our example it will be able to interpret handwriting, but ANN's can in principle be used for all kind of predictions ranging from ecological problems to stock-market fluctuations and weather predictions. It is beyond the scope of this introduction to go into the specific tricky caveats of neural networks and other techniques. However, it is important to realize that there is essentially no difference between these advanced approaches and classical statistical models in the fact that they assume straightforward causality from input to output with all problems of misinterpretation and extrapolation mentioned earlier.

In summary, pragmatic prediction using empirical models may appear attractive, but if we want to predict new developments we have no other option than to base our predictions on a mechanistic insight in the functioning of the system.

Simulation models

In many cases we actually know quite a lot about the main processes that drive a system. Therefore, a logical approach is to combine all that knowledge into a mechanistic computer model, which can then be used to simulate what would happen in different scenarios. For several chemical and physical problems on limited scales this is working relatively well. For instance, the effect of discharge of cooling water of a power plant on the temperature of a river can be computed well. Moreover, the effect of industrial wastewater loaded with various chemicals on the concentration of these chemicals in sediments, water and organisms can also be computed reasonably well. By contrast, effects of CO_2 emissions on the climate and the ocean currents are much more difficult to compute. Although uncertainty is high in such global circulation models, there is no alternative. It is important to have a best guess of our impact on the planet, and we cannot use the empirical approach of trial and error, which still is an option for finding out the effect in small scale situations like lake pollution (simply look at the effect in a test set of 100 comparable lakes).

Predictive simulation models also work surprisingly bad for ecological problems. In the early seventies there was still a great optimism about the possibilities of constructing detailed simulation models for predicting the behavior of ecosystems. Cooperation of groups of experts on all relevant biological and technical sub-topics led to models integrating the available knowledge as much as possible. The model CLEAN (Bloomfield et al., 1974), constructed as part of the International Biological Program research is a good example of this approach. The model contains a diverse spectrum of components like several fish species, algae, zooplankton, aquatic macrophytes, invertebrates and nutrients, formulated in 28 differential equations. The idea of such modeling approaches was that in the course of the modeling process lacking information could be identified, and filled in after additional experimental research. The latter, however, appeared a mission impossible. The number of parameters in those complex models is very large, and the value of many parameters cannot be determined within a reasonable amount of time, if measurable at all.

The common solution is to estimate the remaining parameter values by fitting the model predictions to field data, so-called 'tuning'. A wide array of sophisticated numerical techniques is available for this purpose, and often an impressively good fit is obtained. However, this success is illusive. The problem is that a certain system behavior can often be produced from many different parameter settings. This phenomenon is nicely illustrated by Simons & Lam (1980) who show that even for relatively simple phytoplankton-nutrient models the same patterns can be produced with completely different parameter settings and also from totally different models. Obviously, tuning of complex

ecological models easily leads to good results for the wrong reasons. A good fit does not guarantee any realism of parameter values or model structure. As a consequence, such simulation models have basically the same problems as empirical input-output models. The assumed causal relations underlying the model need not be true, and therefore extrapolation to new situations easily leads to non-sense predictions.

Barriers to a mechanistic understanding of nature
The classic scientific approach

The classical scientific way to unravel how the world works is the so-called 'hypothetico-deductive reasoning'. The main ideas were advocated as early as 1620 by Francis Bacon in his *Novum Organum*, and elaborated later by the influential science philosopher Karl Popper. The steps of this approach which has also be called 'strong inference' are:

1) Devising several alternative hypotheses;

2) Devising a crucial experiment (or several of them), with alternative possible outcomes, each of which will, as nearly as possible, exclude one or more of the hypotheses;

3) Carrying out the experiment so as to get a clean result;

1') Recycling the procedure making sub-hypotheses or sequential hypotheses to refine the possibilities that remain; and so on.

Although this sounds convincing there are some fundamental reasons why this rigid approach is of very limited use in a science dealing with complex systems like ecology. Probably, the most basic argument is that strong inference assumes that the competing hypotheses to explain observed phenomena are general and mutually exclusive, whereas in complex systems in the real world many different entangled causes may contribute to an observed pattern.

Problems of multiple causality

Arbitrarily, we consider the problem of multiple causality here on three different levels (Fig. 1).

In the first place, we have to face that any observed phenomenon can in theory be explained from many different mechanisms. Trivial as it is, this appears hard to keep in mind in every day scientific life. The well-known problem associated with it is that one tends to become attached quickly to ones own tentative hypothesis causing an inevitable bias in further research. A famous treatment of this problem is given in an early paper by Chamberlin (1897) that argues that such affection for ones intellectual child is a grave danger as "love was long since discerned to be blind". "There is then the imminent danger of an unconscious selection and of a magnifying phenomena that fall into harmony with the theory and support it and an unconscious neglect of phenomena that fail of coincidence". In addition to this personal affection problem, there is an interesting intrinsic positive feed-back mechanism in scientific discussion that tends to blow up a hypothesis to a paradigm status: Every time an idea is cited, it gains momentum, becomes more of a "truth" and is more likely to become cited again. A good case study of this problem is described by Elner & Vadas (1990) who analyzed the line of research that started in the early 1970s to find the explanation for a population explosion of sea urchins and the resulting intensive grazing that converted macroalgal beds into barrens along the

Atlantic coast of Nova Scotia. They show how the explanation of this phenomenon as the effect of reduced numbers of lobsters (considered a key-stone predator of the sea urchins) became strongly dominant despite the lack of convincing evidence, and highlight how results against this explanation were circumvented.
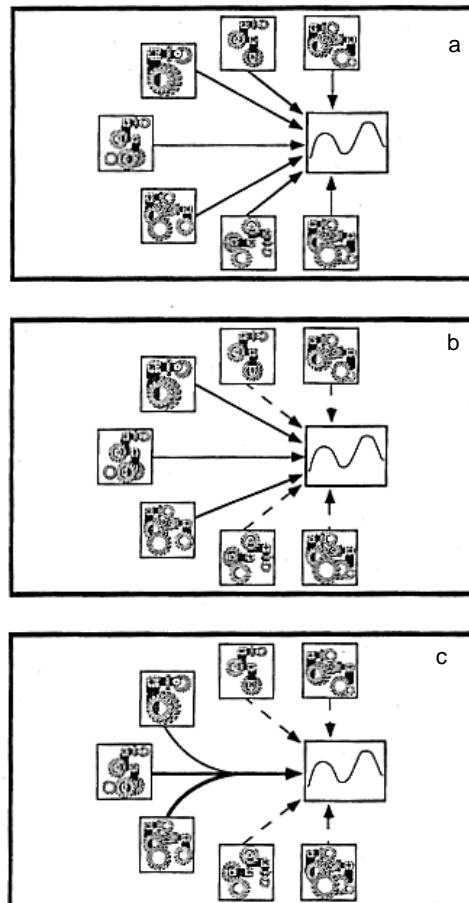
Fig. 1. Three levels of complication in ecological causation:

a) Any observed phenomenon can in theory be explained from many different mechanisms
b) Different mechanisms may actually be responsible for a similar phenomenon in different cases
c) *Several mechanisms may act simultaneously to produce a certain phenomenon that could in theory result from the single ones also.*

A second, equally trivial point to note (Fig. 1) is that the same phenomenon can not only be explained from different theories, but is in fact often caused by different mechanisms in nature in different cases. A man's death can be caused by a bullet, by a fall, by bacteria and in many other ways. This simple complication of causation already presents a difficulty for philosophers who try to define causes and effects in terms of necessary and sufficient conditions (Edwards, 1972). As an ecological example, consider the phenomenon that lakes stay turbid despite a reduction of the

phosphorus loading. In one lake this may be due to wind induced resuspension of the sediment, in a second lake to a release of phosphorus from the accumulated pool in the lake bottom and in a third lake to the impact of the fish community. It is tempting to suggest that a mechanism shown to be responsible for a certain phenomenon in a certain case study is the explanation for this phenomenon in general, and controversies in scientific discussion arise easily if different parties have different case studies in mind.

The above considerations still address only part of the complication of unravelling ecological causation. The full-scale problem comes with the fact that in most real world situations several mechanisms act simultaneously to produce a certain phenomenon that could in theory be explained from each mechanism alone also. Actually, the above-mentioned causes of lake turbidity will practically always co-occur. One of the mechanisms will often dominate, but dominance will differ from case to case and may even shift in time. It is remarkable to note that in the extensive philosophical literature on causation, this common complication is hardly at all touched (Edwards, 1972). Multiple causality makes application of the rigid strong inference method to ecology of little use since, as Quin & Dunham (1983) note, "it is not possible to perform 'critical tests' to distinguish between the 'truth' of processes occurring simultaneously". Multiple causality makes the clean deductive style of reasoning popular among philosophers of science seem futile. The question is not whether a mechanism acts, but rather how important it is in causing an observed phenomenon.

Ockham's Razor (adapted from source ….)
It is important in this context to discuss a famous tool for selecting among alternative theories called *Ockham's razor*. Ockham's Razor is the principle proposed by William of Ockham in the fourteenth century: ``Pluralitas non est ponenda sine neccesitate", which translates as "entities should not be multiplied unnecessarily".

In many cases this is interpreted as "keep it simple", but in reality the Razor has a more subtle and interesting meaning. Suppose that you have two competing theories which describe the same system, if these theories have different predictions than it is a relatively simple matter to find which one is better: one does experiments with the required sensitivity and determines which one gives the most accurate predictions. For example, in Copernicus' theory of the solar system the planets move in circles around the sun, in Kepler's theory they move in ellipses. By measuring carefully the path of the planets it was determined that they move on ellipses, and Copernicus' theory was then replaced by Kepler's.

But there are theories, which have the very same predictions and it is here that the Razor is useful. Consider for example the following two theories aimed at describing the motions of the planets around the sun

- The planets move around the sun in ellipses because there is a force between any of them and the sun, which decreases as the square of the distance.
- The planets move around the sun in ellipses because there is a force between any of them and the sun, which decreases as the square of the distance. This force is generated by the will of some powerful aliens.

Since the force between the planets and the sun determines the motion of the former and both theories posit the same type of force, the predicted motion of the planets will be identical for both theories. The second theory, however, has additional baggage (the will of the aliens), which is unnecessary for the description of the system.

If one accepts the second theory *solely on the basis that it predicts correctly the motion of the planets* one has also accepted the existence of aliens whose will affect the behavior of things, despite the fact that the presence or absence of such beings is irrelevant to planetary motion (the only relevant item is the type of force). In this instance Ockham's Razor would unequivocally reject the second theory. By rejecting this type of additional irrelevant hypotheses guards against the use of solid scientific results (such as the prediction of planetary motion) to justify unrelated statements (such as the existence of the aliens), which may have dramatic consequences. In this case the consequence is that the way planets move, the reason we fall to the ground when we trip, etc. is due to some powerful alien intellect, that this intellect permeates our whole solar system, it is with us even now.., and from here an infinite number of paranoid derivations. For all we know the solar system is permeated by an alien intellect, but the motion of the planets, which can be explained by the simple idea that there is a force between them and the sun, provides no evidence of the aliens' presence nor proves their absence.

A more straightforward application of the Razor is when we are faces with two theories, which have the same predictions, and the available data cannot distinguish between them. In this case the Razor directs us to study in depth the simplest of the theories. Although this sounds reasonable it is very important to realize that this does *not* guarantee that the simplest theory will be correct. Indeed, in biology a tendency of many modelers to adhere too much to the simplest explanation has created quite a tension between these theoreticians and field biologists and experimentalists.

Models as tools for scientific advancement: 'minimal models'
We now turn to a special class of models, which is specially designed to aid in the scientific approach to unravel how complex systems such as organisms, ecosystems and the climate work. These models often referred to as 'minimal models' or 'strategic models' typically focus on one specific mechanism. An important merit is that they can often show an explanation for a certain phenomenon that would not easily be obtained by using just common sense.

Obviously, in view of the reflections on causality presented earlier, even if the model seems to work well, we should be careful not to assume immediately that the modeled mechanism is *the cause* of that phenomenon in any real world situation. In biology, failure to recognize this status of minimal models and their resulting possibilities and limitations for unraveling real world systems has contributed significantly to the antagonism between 'theoreticians' (modelers) and 'naturalists'. Theoreticians, for instance, tend not to pay much attention to the fact that their natural objects of study usually corresponds to only one of many hypotheses for a phenomenon. They rather feel that their job is done when one of these hypotheses is satisfactory modeled. To naturalists this attitude tends to give the impression that the theoretician pretends to give the one and only explanation for the phenomenon in nature. Since it is always obvious that many important factors are left out of the model, this claim may seem so ridiculous that it is not even worth responding to. The situation is easily polarized if theoreticians exaggerate their love for Ockham's razor and consider their explanation for the phenomenon better for the mere reason that it is simpler, i.e. that it requires less biological detail. Probably this will exactly be the reason why many naturalists dislike it, since their ethic is rather one of perceptiveness of the richness of biological detail than of simplicity.

A characteristic critique to a scientific article based on a minimal model could be: "the model contains so few biology that it seems unlikely to learn us anything about real life". A comment remarkably similar to one on experiments saying that "this work on aquaria without sediment is irrelevant to natural systems since all lakes have sediment". The crucial misunderstanding illustrated by these (real) comments is that the experiments or models are intended to capture the whole system whereas they are obviously just focusing on specific mechanisms.

Basically the same misunderstanding tends to hamper the communication between experts and a theoretician trying to cooperate in putting up a minimal model. The expert will typically want to incorporate every aspect that is known or likely to be of influence, whereas the theoretician tries to keep out as much as possible. It may seem unlikely that at the end of this process of seemingly arbitrary choices a model will come out that has exactly the components that it "needs". However, again there is a misinterpretation of the intentions in this. If the model would be intended to give a quantitatively good description of the systems behavior, every important aspect should indeed be included, but unlike in the case of predictive simulation models, with a minimal model we are in the comfortable situation that there is no need to be exhaustive. We still want to make clever choices, so as to address mechanisms that are likely to be important in nature. However, for the soundness of scientific argumentation it actually does not matter what we put in and what we leave out, as long as the results are interpreted in a proper way.

Proper interpretation requires that we realize that a model does nothing else than showing the result of the incorporated processes and relationships. If we are lucky these results are at first sight counter-intuitive. We then obtained a new explanatory mechanism on the checklist (Fig. 1) that we can use in the search for the real causes of a phenomenon in nature. A mechanism that because of its counter-intuitiveness would probably never have occurred to us if we would only have used common sense but that is just as likely to occur in nature. A somewhat counter-intuitive explanation obtained from a model may feel less safe and convincing than a common sense biologist's explanation, but the reliability of human feeling in these matters is at least questionable. After all, common sense is also the thing that tells us that the world is flat.

In addition to suggesting alternative explanations for phenomena observed in nature, minimal models can effectively enhance our insight in mechanisms that are difficult to grasp intuitively. Playing around with such models gives a feeling for the effects of the intertwined feedback systems often encountered in natural systems. Obviously, the question whether the hypothesized mechanisms are actually responsible for the real world behavior can only be answered by studying the real system. However, the search images needed to detect intricate regulatory mechanisms can effectively be obtained by experimenting with models.

Note that there is not really an essential difference between the small minimal models and the larger elaborate simulation models used for prediction. Both are usually sets of differential equations representing the dynamics of important components of the system. Indeed, experiments with large

simulation models can also generate valuable insights in how the real system may work. However, small models are simply more suitable for that purpose because they are easier to understand.

Modeling in practice: formulation, analysis and interpretation

Once you get the hang of it, modeling is great fun. There is some magic to the possibility to build a kind of toy-world, which starts a life on its own in your computer. As a result it is sometimes difficult to remember that modeling is a tool, not a goal. Any modeling exercise should obviously start with a question. What is it we want? Explanation for a certain puzzling phenomenon? Prediction of the impact of pollution? Subsequently, we should ask ourselves two questions:

1) Is a model the best way to solve this question. And if so,
2) Which type of model would be most appropriate.

Formulation: making a model

In case this leads to the conclusion that a mechanistic mathematical model would be of help, the process of identifying important mechanisms starts. In case of a minimal model the choice is somehow arbitrary as long as we are careful in our final conclusions. In case of a predictive model, it is very important to have good quantitative information about all the important processes. The next step is to represent the processes in a mathematical form. Usually, standard formulations that have been used before for certain processes are the best choice, but sometimes one can defend new mathematical ways of describing certain processes.

Analysis: finding out how the model behaves

Some properties of simple sets of equations (especially minimal models) can be analyzed analytically with pen and paper. However, mostly we will want to use the computer to help analyzing the 'emergent properties' of the model. Many different techniques and software packages are available for this step.

Interpretation: what we can learn from this about the real world

The last step is to formulate what the results of the modeling exercise tell us about the initial real world problem. This is probably the most difficult step, and the step, which is often taken in a much too sloppy way in practice. It will be clear after reading this introduction, that it is equally important to stress what can be concluded from a model, as to stress what cannot.

Assignments
- The incidence of heart attacks and other cardio-vascular problems is correlated to the cholesterol content of the blood of people. This has led to a tendency to eat food, which

contains little cholesterol. Which causalities are implicitly assumed here, and how could we test these.

- Analysis of ancient ice-layers and other paleaeo evidence has revealed that the CO_2 content of the atmosphere was systematically low during ice ages and high during warm periods in between. Ice ages are invoked by reduced incoming solar energy due to variations in the orbit of the earth. What does this suggest about the causal relationship between CO_2 content of the atmosphere and temperature? What would be the use of a regression model of the relationship between these parameters based on the reconstructed time series of ancient earth history? Draw a schedule of the relationships including human emissions. What kind of models could be of help to elucidate our role? Which processes should be included?

- Mobile telephones emit radiation which some consider dangerous for the users. Which kind of model (or models) would you advice to explore whether there are reasons for such fear?

- Give an example of a correct use and of an incorrect use of Ockham's razor.

References

Bloomfield, J.A., Park, R.A., Scavia, D. and Zahorcak, C.S., 1974. Aquatic modelling in the eastern deciduous forest biome, U.S.- International biological program. In: E. Middlebrooks, D.H. Falkenberg and T.E. Maloney (Editors), **Modeling the Eutrophication Process**. Ann Arbor Science, Ann Arbor, MI, pp. 139-158.

Chamberlin, T.C., 1897. The method of multiple working hypotheses. Journal of Geology, 5: 837-848.

Edwards, P., 1972. The encyclopedia of philosophy. Collier MaxMillan Publishers, New York.

Elner, R.W. and Vadas, R.L., 1990. Inference in ecology the sea urchin phenomenon in the northwestern Atlantic ocean. Amer.Naturalist, 136: 108-125.

Montgomery, D.C. and Peck, E.A., 1982. Introduction to Linear Regression Analysis. Wiley, New York, 1-504 pp.

Quinn, J.F. and Dunham, A.E., 1983. On hypothesis testing in ecology and evolution. Amer.Naturalist, 122: 602-617.

Simons, T.J. and Lam, D.C., 1980. Some limitations of water quality models for large lakes a case study of Lake Ontario USA Canada. Water Resources Research, 16: 105-116.

Vollenweider, R.A., 1977. Advances in defining critical loading levels for phosphorus in lake eutrophication. Memorie dell'Istituto Italiano di Idrobiologia: 53-83.

Equilibrium and Bifurcations

(*based on Scheffer, 2009*)

**Tipping points and resilience**

Suppose that you are in a canoe and gradually lean farther and farther over to one side to look at something interesting underwater. Leaning over too far may cause you to capsize and end up in an alternative stable state upside down. Although the details of the theory of alternative stable states may appear tricky, several key properties can be seen in this simple example. For instance, returning from the capsized state requires more than just leaning a bit less to the side. It is difficult to see the tipping point coming, as the position of the boat may change relatively little up until the critical point. Also, close to the tipping point, resilience of the upright position is small, and minor disturbances such as a small wave can tip the balance.

Resilience, defined as the ability of a system to recover to the original state upon a disturbance, is a tricky issue. The famous seventeenth-century Swedish ship the Wasa is an example in case. Named for the royal house of Wasa, the ship was built as the most prestigious ship of the navy of Gustavus Adolphus to combat the Polish on the Baltic. Against the advice of his engineers, the king insisted on an extra layer making the ship higher. The result looked impressive. However, when the ship set sail on its maiden voyage, a sudden breeze of wind made it sink within minutes, not a mile from land. In a similar vein, we are reminded almost every year of the instability of ships by tragic accidents with ferries. Overloading is an especially treacherous aspect, as it can reduce resilience of the correct position in an unperceived way. A relatively small disturbance like wave action or too many people moving to one side might then tip the boat.

In such everyday examples of systems with alternative stable states, the consequences are intuitively straightforward. Nevertheless, with respect to complex systems such as societies and ecosystems, the idea is perhaps somewhat counterintuitive. Could coral reefs, shallow lakes, nerve cells, the climate, or public attitude really tip over like a canoe? If so, can we manage or predict such shifts? In this chapter we will go through the basic theory and explain why some biological systems might have tipping points.

Equilibrium in dynamical systems

The theory of dynamical systems is a branch of mathematics used to describe the interaction between all kinds of state variables. That sounds fairly broad, and indeed it is. A cell, a fish, a population, or the Earth may be considered a dynamical system - that is, their state can be understood as the result of underlying governing processes. For instance, the temperature of the Earth is largely the result of heat gained from solar radiation and heat lost by radiation of heat back into the atmosphere. The result is a **stable equilibrium**. A state is called an equilibrium when it holds that if you start exactly in that
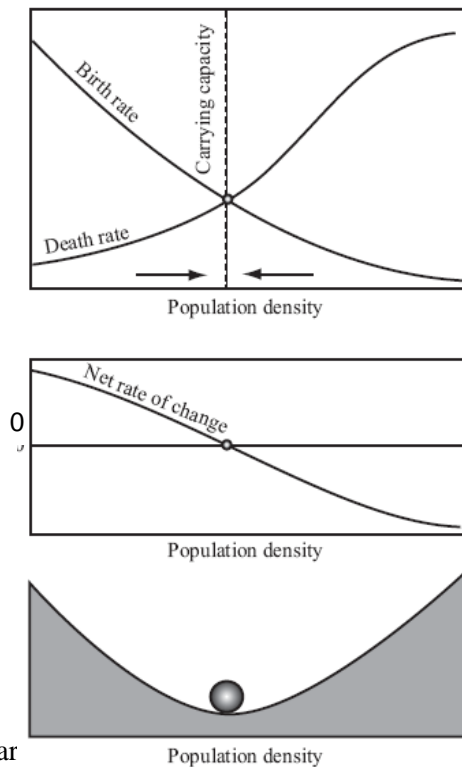
Figure 2P.1. The concept of a stable dynamic equilibrium illustrated for the case of a hypothetical population that settles at a density that corresponds to the carrying capacity of the environment (top panel). The slope of the hills in the stability landscape (bottom panel) corresponds in a negative sense to the rate of change in the population density (middle panel), which is the net result of per capita birth and death rates (top panel).

par Each equilibrium can be (locally) stable or unstable and stability of an equilibrium means in the case of the temperature of the Earth that if the temperature were brought into a slightly different state, it would move back to the original equilibrium value. For instance, if the Earth were cooler than the equilibrium temperature, it would radiate less into the atmosphere. Since the incoming radiation remains the same, this means that the Earth would warm up, thus moving in the direction of the equilibrium temperature. On the other hand, if the Earth were warmer, it would radiate more, and thus lose heat and cool down to the equilibrium again.

Another well-known example is that of a population that has reached the **carrying capacity** of the environment (Figure 2P.1). Such a population is at an equilibrium density resulting from a balance between birth and death rates. Again, it is a stable equilibrium. If a proportion of the population were wiped out by an adverse advent, there would be more resources for the survivors. This promotes birth rates and reduces death rates so that the population grows back to the equilibrium density. On the other hand, if densities exceed carrying capacity, reduced birth and increased mortality will push it back to the equilibrium. The overall rate of change in the population is the net result of gains and losses; in equilibrium, it is zero.

An attractive way to depict the stability of a system is by means of a **stability landscape** (Figure 2P.1, bottom panel). The slope at any point in such a landscape corresponds in a negative sense to the rate of change. Thus, at equilibrium, where the rate of change is zero, the slope is also zero. You can now imagine the system like a ball settling in the lowest point, representing the equilibrium. This "**ball in a cup**" analogue should not be taken too literally, for instance the ball usually will slide towards the equilibrium and not overshoot like a real ball would do. Nonetheless, it is a good intuitive aid for grasping the essence of the matter.

For obvious reasons, a stable state as depicted in these graphs is also called an **attractor**. The idea of an equilibrium is an important concept, but it is of course so that many systems are in reality not in equilibrium. The temperature of the Earth or the densities of populations are never constant. Fluctuations in the environment and all kinds of smaller or larger perturbations prevent the systems from settling into an equilibrium. Obviously, the equilibrium state can change if conditions are altered. For instance, as fossil fuel combustion causes greenhouse gas levels to increase, the Earth will retain relatively more heat and become warmer. There can be also other kinds of attractors, such as cycles or strange attractors, that prevent systems from settling to a stable state, even if they were in a perfectly constant environment. Moreover, systems can have several alternative attractors, separated by repelling (=unstable) points, cycles, or other structures, and **critical transitions** happen if a system shifts from one attractor to another.
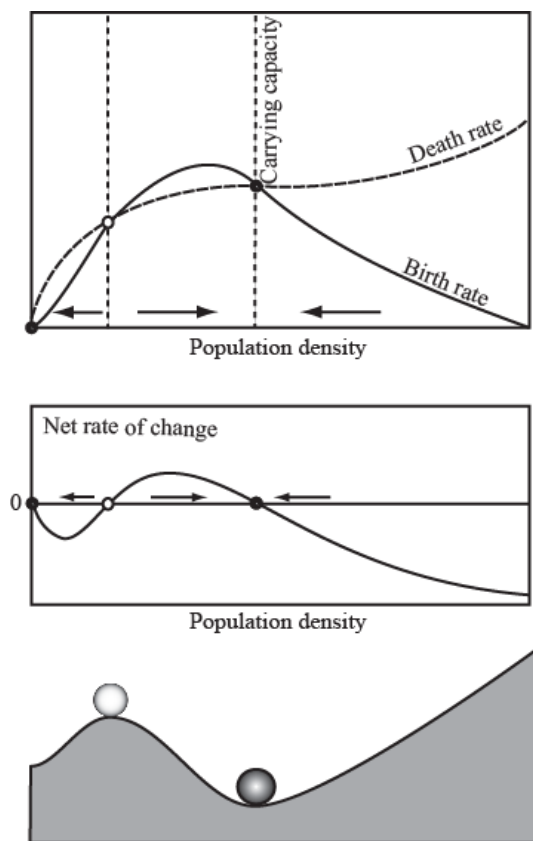
Figure 2P.2. The concept of alternative equilibria illustrated for the case of a hypothetical population that experiences an Allee effect, implying that at low densities, the net growth of the population is negative because of increased per capita mortality and reduced per capita birth rates. As a result, a situation in which the population density is zero represents a stable state. Only if the critical density that marks the border of the basin of attraction of the nil state (small open circle in upper panel and light-colored ball in lower panel) is exceeded the population grows and ends up in the alternative stable state at carrying capacity.

Alternative equilibria

As a simple example of a system with alternative attractors, imagine a population of animals that runs into trouble if its numbers get too low. This may happen, for instance, if finding a mate becomes too difficult at low densities. Also, this can occur in species that congregate to protect themselves against predators. In such cases, mortality may go up at low population densities, but birth rates may fall too. For example, animals such as flamingos and penguins do not breed unless they are surrounded by many other mating individuals of the same species. The first to suggest that population growth can become depressed at low densities was the American zoologist Warder Allee, and the phenomenon is now commonly referred to as the **Allee effect**. Clearly, this is a highly important mechanism when it comes to understanding the extinction of endangered species. If the Allee effect is strong enough, it implies that a population can go into free fall if its density goes below a certain critical level (Figure 2P.2). In that case, the population has two alternative stable states: one at carrying capacity as in the case described earlier; the other at density zero. In the previous simple growth model (Figure 2P.1), a zero density is also an equilibrium state, as the absence of parents results simply in no offspring. In that case, however, a small addition of animals will be enough to kick off a population increase that stops only when the carrying capacity is reached, therefore it is an unstable equilibrium. In contrast, a population with a strong Allee effect is trapped into the zero-equilibrium state. A small initial population number is drawn back into the zero-state, as long as it is below the critical density (open dots in Figure 2P.2). As you will see later, plant populations can also have such a critical density. This happens especially under harsh conditions where a critical plant density is needed in order to

"**engineer**" the environment sufficiently to make it suitable for plant growth. A simple mathematical model of a population with an Allee effect can be found in Box 2P.1.

In the following, we will look more closely at the general consequences for the dynamics of systems.

BOX 2P.1: The Allee effect

Some populations have a threshold density below which they go into free fall towards extinction. In the main text, different mechanisms that can cause this phenomenon are explained. One of the many possible equations that can describe such an Allee effect in a single population is:

$$\frac{dN}{dt}(t) = rN(t)\left(1 - \frac{N(t)}{K}\right)\left(\frac{N(t) - C}{K}\right)$$

In which: N(t) is the population size at time t (g m–2), r is the relative growth rate (d–1), K is the carrying capacity (g m–2) and C a critical density (g m–2). The difference from the classic logistic growth equation is the multiplication by the last term $\left[\frac{N-C}{K}\right]$, which implies that the population growth becomes zero not only at extinction $(N = 0)$ and carrying capacity $(N = K)$, but also at the Allee extinction threshold (C). Below this critical density $(N = C)$, net growth is negative and the population goes extinct. Parameters can be chosen on the basis of data, but obviously the extinction threshold (C) should be larger than zero to obtain the typical Allee effect (and by definition C < K).
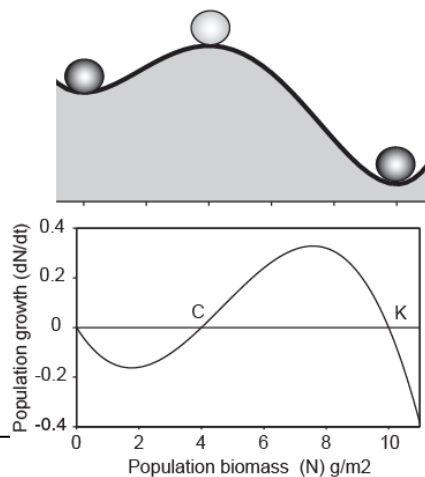

Figure 2P.3.(a) The stability landscape. (b) The population growth of a logistically growing species with an Allee effect ($r = 1$, $C = 4$ and $K = 10$). It can easily be seen that this model has three equilibria, of which there are two stable and one unstable (Which and why?).

Catastrophic shifts and hysteresis

The existence of **multiple stable states** has profound implications for the way in which a system responds to changing conditions. Mostly, the equilibrium of a dynamical system changes smoothly in response to changes in the environment (Figure 2P.4a). Also, it is quite common that the system is rather insensitive over certain ranges of the external conditions, while responding relatively strongly around some threshold condition (Figure 2P.4b). For instance, mortality of a species usually increases sharply around some critical concentration of a toxicant. In such a situation, a strong response occurs

when a threshold is passed. Such thresholds are obviously important to understand. However, a very different, much more extreme kind of threshold than this occurs if the system has alternative stable states. In that case, the curve that describes the response of the equilibrium to environmental conditions is typically "folded" (Figure 2P.4c). Note that such a **catastrophe fold** implies that indeed for a certain range of environmental conditions, the system has two alternative stable states, separated by an **unstable equilibrium** (dashed line) that marks the border between the basins of attraction of
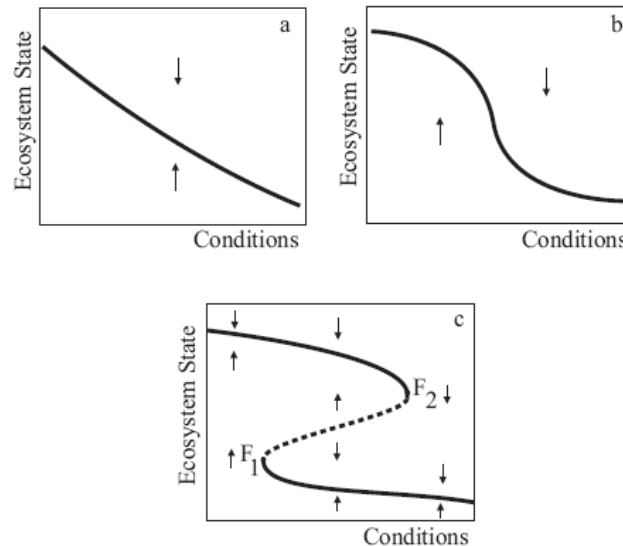
Figure 2P.4. Schematic representation of possible ways in which the equilibrium state of a system can vary with conditions such as nutrient loading, exploitation, or temperature rise. In panels (a) and (b), only one equilibrium exists for each condition. However, if the equilibrium curve is folded backward as in panel (c), three equilibria can exist for a given condition. The arrows in the graphs indicate the direction in which the system moves if it is not in equilibrium (that is, not on the curve). It can be seen from these arrows that all curves represent stable equilibria, except for the unstable equilibrium values on the dashed middle section in panel (c). If the system is pushed away a little bit from this part of the curve, it will move further away instead of returning. Hence, these unstable equilibria represent the border between the basins

the alternative stable states, just as in the example of the population with an Allee effect (Figure 2P.2).

This situation lies at the root of true critical transitions: when the system is in a state on the upper branch of the folded curve, it cannot pass to the lower branch smoothly. Instead, when conditions change sufficiently to pass the threshold (F2, Figure 2P.4), a **"catastrophic" transition** to the lower branch occurs (Figure 2P.4c). Clearly, this is a very special point. In the exotic jargon of dynamical systems theory, it is called a **bifurcation point**. There exist several different kinds of bifurcation points that all mark thresholds at which the system's qualitative behavior changes. For instance, the system may start oscillating, or a species may go extinct at a bifurcation point.

The point (F_2) shown in Figure 2P.4c and 2P.5 marks a so-called **catastrophic bifurcation**. Such bifurcations are characterized by the fact that an infinitesimally small change in a control parameter (reflecting, for instance, the temperature) can invoke a large change in the state of the system if it crosses the bifurcation. While all kinds of bifurcations correspond in a sense to critical transitions, catastrophic bifurcations are really the mathematical analogue of the dramatic transitions, e.g. from

clear water to algal dominated water, that we are trying to understand. The bifurcation points in a catastrophe fold (F1 and F2) are known as **fold bifurcations**.
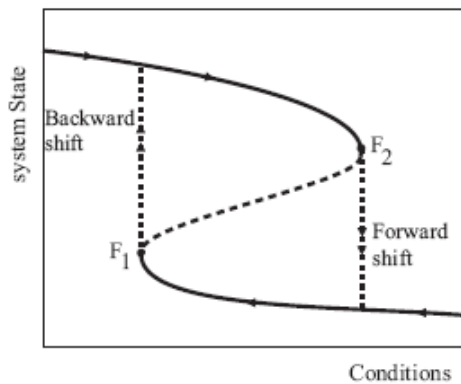
Figure 2P.5. If a system has alternative stable states, critical transitions and hysteresis may occur. If the system is on the upper branch, but close to the bifurcation point F2, a slight incremental change in conditions may bring it beyond the bifurcation point and induce a critical transition (or catastrophic shift) to the lower alternative stable state (forward shift). If one tries to restore the state on the upper branch by means of reversing the conditions, the system shows hysteresis. A backward shift occurs only if conditions are reversed far enough to reach the other bifurcation point F1.

The fact that a tiny change in conditions can cause a major shift is not the only aspect that sets systems with alternative attractors apart from the "normal" ones. Another important feature is the fact that in order to induce a switch back to the upper branch, it is not sufficient to restore the environmental conditions from before the collapse (*F2*). Instead, one needs to go back further, beyond the other switch point (*F1*), where the system recovers by shifting back to the upper branch. This pattern in which the forward and backward switches occur at different critical conditions (Figure 2P.5) is known as **hysteresis**. From a practical point of view, hysteresis is important, as it implies that
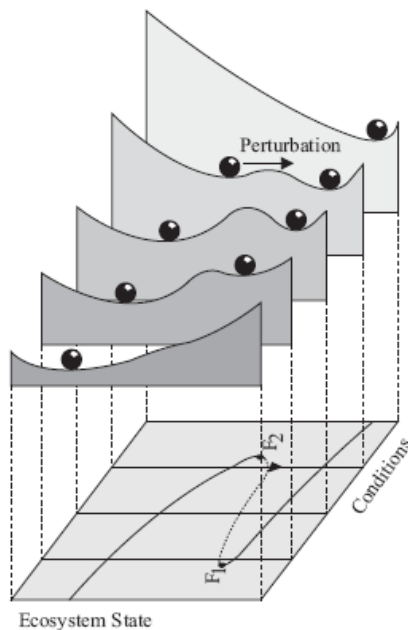
Figure 2P.6. External conditions affect the resilience of multiple stable systems to perturbation. The bottom plane shows the equilibrium curve as in Figure 2.5. The stability landscapes depict the equilibria and their basins of attraction at five different conditions. Stable equilibria correspond to valleys; the unstable middle section of the folded equilibrium curve corresponds to hilltops. If the size of the basin of attraction is small, resilience is small, and even a moderate perturbation may bring the system into the alternative basin of attraction.

18

this kind of catastrophic transition is not so easy to reverse.

The idea of catastrophic transitions and hysteresis can be nicely illustrated by stability landscapes. To illustrate how stability is affected by changes in conditions, we create stability landscapes for different values of the conditioning factor (Figure 2P.6). For conditions at which there is only one stable state, the landscape has only one valley, just as in the case discussed at the beginning of this section (Figure 2P.1). However, for the range of conditions where two alternative stable equilibrium states exist, the situation becomes more interesting. The stable equilibrium states occur as valleys, separated by a hilltop. This hilltop is also an equilibrium (the slope of the landscape is zero). However, this equilibrium is unstable. It is a **repellor**. Even the slightest change away from it will lead to a self-propagating runaway process moving the system toward an attractor.

To comprehend the catastrophic transitions and hysteresis, imagine what happens if you start in the situation of the landscape in front. The system is in the only existing equilibrium. There is no other attractor, and therefore this state is said to be globally stable. Now suppose that conditions change gradually, so that the stability landscape changes to the second or third one in the row. Now there is an alternative attractor, implying that the state in which the system was has become locally (rather than globally) stable. However, as long as no major **perturbation** occurs, the system will not move to this alternative attractor. In fact, nothing would reveal the fundamental changes in the stability landscape. If conditions change even more, the basin of attraction around the equilibrium in which the system rests becomes very small (fourth stability landscape) and eventually disappears (last landscape), implying an inevitable catastrophic transition to the alternative equilibrium state. Now, if conditions are restored to previous levels, the system will not automatically shift back. Instead, it shows **hysteresis**. If no large perturbations occur, it will remain in the new state until the conditions are reversed beyond those of the second landscape.

Resilience as the width of a basin of attraction

In reality, conditions are never constant. Accidental, or stochastic, events such as weather extremes, fires, or pest outbreaks can cause fluctuations in the conditioning factors but may also affect the state directly - for instance, by wiping out parts of populations. If there is only one basin of attraction, the system settles back to essentially the same state after such events. However, if there are alternative stable states, a sufficiently severe perturbation may bring the system into the basin of attraction of another equilibrium state. Obviously, the likelihood that this happens depends not only on the perturbation, but also on the size of the attraction basin. In terms of stability landscapes (Figure 2P.6, if the valley is small, a small perturbation may be enough to displace the ball far enough to push it over the hilltop, resulting in a shift to the alternative stable state. Following Holling (1973), we use the term **resilience** to refer to the size of the valley or basin of attraction around a state that corresponds to the maximum perturbation that can be taken without causing a shift to an alternative stable state.

A crucially important phenomenon in systems with multiple stable states is that gradually changing conditions may have little effect on the state of the system, but nevertheless reduce the size of the attraction basin (Figure 2P.6). This loss of resilience makes the system more fragile in the sense that it can be easily tipped into a contrasting state by stochastic events (i.e. events that occur by chance such

as weather events; see also week 4). This is also one of the most counterintuitive aspects. Whenever a large transition occurs, the cause is usually sought in events that might have caused it: the collapse of some ancient cultures may have been caused by droughts. An intractable conflict may be due to the act of an evil leader. A lake may have been pushed to a turbid state by a hurricane, and a meteor is thought to have wiped out the dinosaurs, leading to the rise of mammals. The idea that systems can become fragile in an invisible way because of gradual trends in climate, pollution, land cover, poverty, or exploitation pressure may seem counterintuitive. However, intuition can be a bad guide, and this is precisely where good and transparent systems theory can become useful. Currently, such a **resilience-based management** style is starting to be used in some systems such as the management of coral reefs, but it requires a major paradigm shift in many other areas.

Phase plane analysis

Here we will introduce the analysis of more complex models. If you want to describe the interactions between for instance two competing species, you need to describe the change in each of the two species. Therefore it is needed to describe the change in two **state variables**: one describing the change in species *A* and one for species *B*. In the particular case of competition between two species, alternative stable states may also arise. The general rule here is that competition can lead to alternative stable states if it is better to have individuals of the same species around than individuals of the other species (i.e. i**ntraspecific competition** is less severe than **interspecific competition**.). This implies a positive feedback in developing towards a monoculture of each of both species (the two alternative attractors). In the following we will show how a phase plane analysis can help to analyze the classical two species Lotka-Volterra competition model.

The idea of a phase plane analysis is to draw the line at which one species is in equilibrium (the **zero-growth isoclines** or **nullclines**) as a function of the other species and vice versa (Figure 2P.7). Such nullclines can be produced, for instance, from the famous Lotka-Volterra competition model:

$$\frac{\mathrm{d}A}{\mathrm{d}t}(t) = r_{\mathrm{A}} A(t) \left(1 - \frac{A(t) + \alpha_{\mathrm{AB}} B(t)}{K_{\mathrm{A}}} \right)$$

$$\frac{\mathrm{d}B}{\mathrm{d}t}(t) = r_{\mathrm{B}} B(t) \left(1 - \frac{B(t) + \alpha_{\mathrm{BA}} A(t)}{K_{\mathrm{B}}} \right)$$

Note that the first part of the growth equation of species *A* is simply the logistic equation (with growth rate r_{A} and carrying capacity K_{A}). However, in addition to the intraspecific competition, that was already included in the logistic equation, species *A* now suffers from competition by the other species $(\alpha_{\mathrm{AB}} B(t))$. The equation for the growth of species *B* is analogous to that for species *A*. The nullclines plotted in Figure 2P.7 and Figure 2P.8 are solutions of the growth equations $\frac{dA}{dt}(t) = 0$ and $\frac{dB}{dt}(t) = 0$. It can be seen from the nullclines that the model assumes that the equilibrium density of each species declines linearly with increasing density of its competitor (Figure 2P.8). Note however that if species *A* is in equilibrium, species *B* might still change, therefore the total model has only an equilibrium if *both* nullclines intersect.
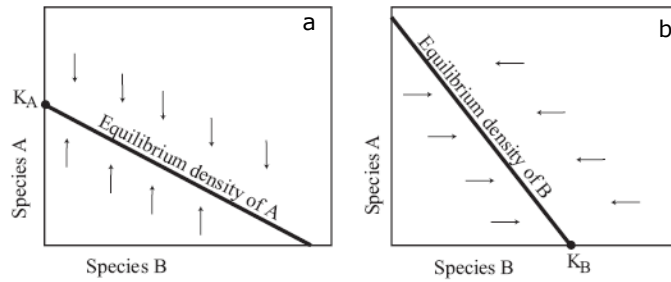
Figure 2P.7. Competition of two species can be seen as a process in which they affect the carrying capacity of each other. In the left panel (a) the nullcline of species A is drawn and on this line the abundance of species B is the only one that changes. Increasing abundance of species B reduces the carrying capacity of A, and hence the equilibrium density to which species A will settle. Moreover, in this plane the change in abundance of species B is not given explicitly. Similarly, in the right panel (b) the nullcline of species B is drawn and on this line the abundance of species A is the only one that changes: the equilibrium density to which species B will settle depends on the density of species A. In this plane the change in abundance of species A is not given explicitly (b). The arrows indicate the direction of change of only one species at a time (so both figures should be combined to see the real change in both species together.

The steepness of the decline depends of course on the intensity of competition. If the competitor has a somewhat different niche, it may be that its suppressing effect is minor. To see the effect of the competition, one should plot the equilibrium lines of the two species
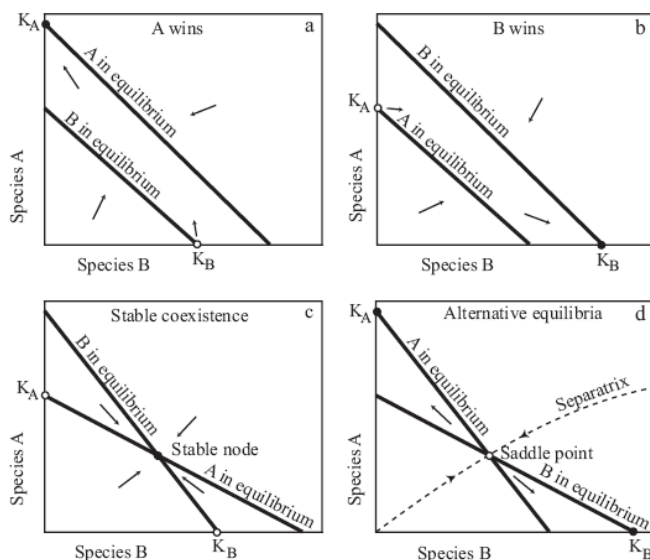

Figure 2P.8. Equilibrium lines for competing species A and B combined reveal the equilibria resulting from the competition. Dots are stable equilibria, whereas open circles represent unstable equilibria.

21

together (Figure 2P.8). If one equilibrium line is entirely above the other, the species with the highest equilibrium wins. If the two lines intersect, this is a point at which both species are at equilibrium. There are two possibilities now, either the intersection point is stable and both species can coexists (Figure 2P.8c), or the intersection point is an unstable equilibrium and there are two alternative equilibrium states, either with species A or with species B. The unstable equilibrium is called a **saddle point** (Figure 2P.8d). Though a saddle point is unstable, it is also attracting from two directions. If you start in Figure 2P.8d with a combination of both species close to the dotted line (the separatrix), you first move in the direction of the saddle point before you bend to either dominance of species A or species B. It can be seen from Figure 2P.8c and d that stable coexistence requires that interspecific competition is relatively weak (i.e. the slope of the nullclines is such that the competitor has less effect on the equilibrium of either species - the slope of both is less steep compared to the line between K_A and K_B), whereas alternative equilibria arise if interspecific competition is stronger than intraspecific competition (implying that the decline in equilibrium density as a function of the density of the competitor is steep). This can also be seen from constructing the resulting arrows from the separate nullclines (Figure 2P.7). Note that above the nullcline of A the biomass of A decreases, below A increases, and right of the nullcline of B its population decreases and left of the B nullcline, B increases.

To see the difference in response to environmental conditions, imagine what happens if the environment changes in a way that is good for species A but has no direct effect on species B. This implies that the equilibrium line for species A in the diagram moves up (as K_A increases), whereas the line for B remains unaltered. If we start in a situation in which species B would win (panel b), this will eventually turn the system into a situation in which A will win (panel a). However, the transition depends on whether coexistence is stable or unstable (Figure 2P.9). In the case of stable coexistence (Figure 2P.8c; Figure 2P.9a), as soon as the lines cross, the resulting stable point at which both species are in equilibrium will move gradually to the upper left, implying that species B becomes gradually less abundant in the coexistence equilibrium. When the intersection point eventually reaches the vertical axis, species B goes extinct, and a monoculture of A remains. This is a so-called **transcritical bifurcation**. It is not a catastrophic bifurcation, as the change for species B from rare to extinct is only a small (albeit important) one.

Now consider what happens if the same environmental change occurs in a system with alternative equilibria (Figure 2P.8d; Figure 2P.9b). Suppose that the system is in the equilibrium dominated by species B (at K_B) while species A is trying to invade from elsewhere. If the lines cross, a further increase in K_A will cause the unstable intersection point to slide to the lower right. However, the B monoculture equilibrium (K_B) remains stable as long as the unstable intersection point does not hit it. Nonetheless, it can be seen that the resilience of this monoculture decreases and species B becomes vulnerable to invasions of species A. As the saddle moves toward the B monoculture point, the attraction basin of this equilibrium thus shrinks. This implies that moderate invasions of species A may be enough to trigger a shift to the A monoculture state. Eventually, if the equilibrium line of A moves enough to let the saddle point collide with the B monoculture equilibrium (K_B), this monoculture becomes unstable. Even the slightest invasion by A will now cause a runaway shift to a monoculture of A at K_A. Note that as long as A remains absent, the shift away from the B monoculture will not occur. Such an unstable point in which one variable is zero is sometimes called a **trivial**

equilibrium. This is probably a common situation in nature, as illustrated by the fact that biological invasions are frequently successful. Many species are not absent because they cannot coexist with the rest of the community, but rather because they never arrived.
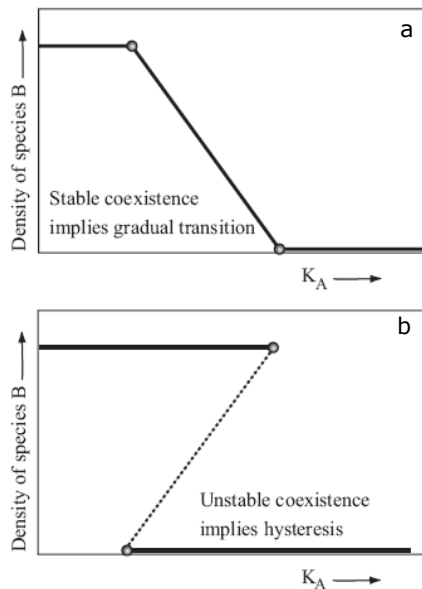
Figure 2P.9. In competition, environmental change that favors species A but has no direct effect on species B can lead to a gradual (a) or catastrophic (b) transition, depending on the relative strength of interspecific versus intraspecific competition (see text for explanation).

References

Holling, C. S., 1973. Resilience and stability of ecological systems. Annu. Rev. Ecol. Syst. 4: 1-23.

Scheffer, M., 2009, Critical Transitions in Nature and Society. Princeton University Press, Princeton and Oxford, 400 pp.

Aspects of equilibria in 1 dimensional systems

Equilibria in systems and models have many interesting characteristics, as described in the lecture Equilibrium and Bifurcations. In the following examples we will explain some of these characteristics in more detail, such as stability and bifurcation using a number of different models as example.

Stability analysis in exponential growth and decay   The exponential growth or decay model is one of the most simple models to describe for instance bacterial growth ($r > 0$), or the decay of a substance ($r < 0$): $dx(t)/dt = rx(t)$. The equation states that the rate of change of the state ($dx(t)/dt$) is linearly related to the amount of x present. We can obtain the analytical solution by separation of variables. The solution we obtain with the initial value $x(0) = x_0$ is

$$x(t) = x(0)e^{rt}$$

The dimensions in the text and figures for x is mass per unit area, e.g. with units g m^{-2}, and time, T, e.g. with the unit of day.

Since an equilibrium is the situation where the system is not changing anymore, the only possibility for this system to be in an equilibrium is if $x = x_{eq.} = 0$; so no bacteria present or no substance left anymore.

We can graphically determine whether this equilibrium is stable by plotting the derivative of x ($dx(t)/dt$) or, for brevity, (x') as a function of x.

For exponential growth ($r > 0$) this is depicted in Figure 2T.3.

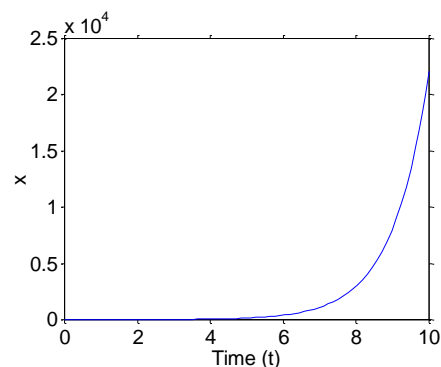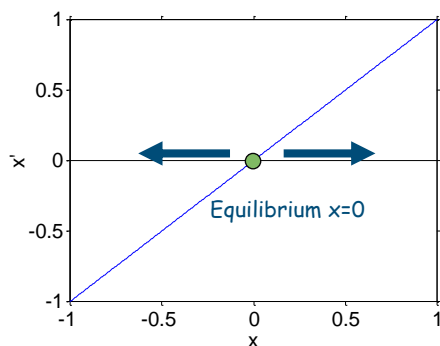$r > 0$ Unstable: grows rapidly to infinity

Figure 2T.3-a. Rate of change $(x' = dx(t)/dt)$ *Figure 2T.3-b. State x plotted against time.*
  plotted against state x.

The left panel shows that the only value for which x' equals zero is $x = 0$. If the state x is positive then $x' > 0$ and thus x is increasing, and will go to the right. If x is negative then $x' < 0$ and thus x is decreasing and going to the left. As the arrows point away from the equilibrium, $x = 0$ is an unstable equilibrium. Note that in reality x cannot be negative. However, for the mathematical analysis this is no problem. The right panel shows the time evolution of x for the case of a positive initial value of x and a positive r. As can be seen x increases for ever.

For exponential decay $(r < 0)$ the graphical analysis is depicted in Figure 2T.4.

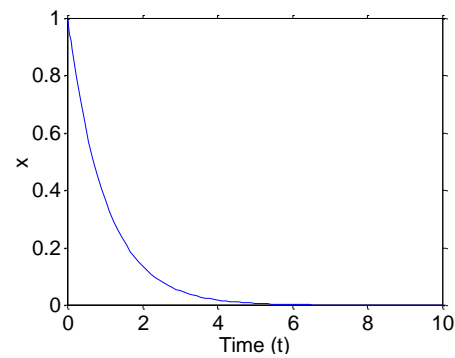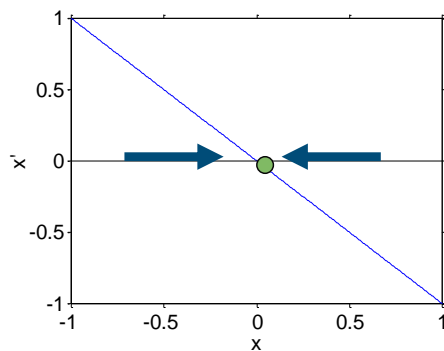$r < 0$ 　　　　　　　　　　　　　　　　　Stable: approaches zero, (initial condition $= 1$)

Figure 2T.4-a. Rate of change $(x' = \mathrm{d}x(t)/\mathrm{d}t)$ 　 *Figure 2T.4-b. State x plotted against time.*
plotted against state x.

Also here the left panel shows that the only value for which x' equals zero is $x = 0$. If the state x is positive then $x' < 0$ and thus x is decreasing, and will go to the left. If x is negative then $x' > 0$ and thus x is increasing and will go to the right. Therefore, as the arrows point towards the equilibrium, $x = 0$ is a stable equilibrium. Again x cannot be negative in reality, but for the mathematical analysis this is no problem. The right panel shows the time evolution of x for the case of a positive initial value of x and a negative r. As you can see, x goes to zero.

This exponential model is very simple and especially for growth $(r > 0)$ not very realistic. Only for a limited time period it may describe growth in natural populations. For longer periods, however, the growth will soon become limited, for example with respect to the food needed for growth or the amount of space available for a population of animals. Generally, this is called the carrying capacity of the system. In the classical logistic growth model such limitations are introduced by the term $1 - (x(t)/K)$ that is added to the exponential growth equation to give

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = rx(t)\left(1 - \frac{x(t)}{K}\right) \tag{2T.3}$$

where K is the carrying capacity. The equilibrium in the logistic equation will be reached if the rate of change of the state is zero, so either at $x(t) = 0$ or $x(t) = K$.

Again we can graphically determine the stability of this equilibrium by plotting the derivative of x (x') as a function of x, Figure 2T.5. Note that the graph is a parabola.
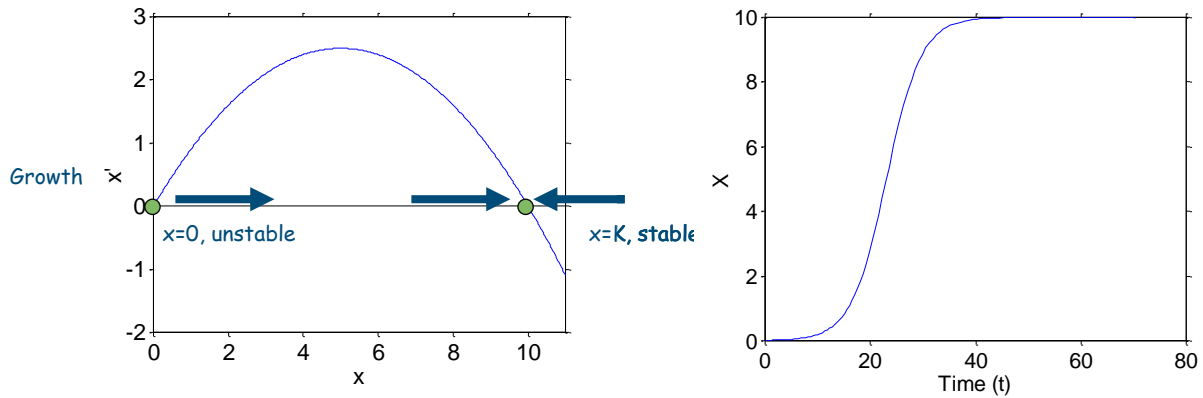

Figure 2T.5-a. Rate of change $(x' = \mathrm{d}x(t)/\mathrm{d}t)$ plotted against state x.

Figure 2T.5-b. State x plotted against time, using initial condition $x(0) = 0.01$.

The graph in Figure 2T.5 is plotted for $r > 0$ and $K = 10$. There indeed appear to be two values for which x' equals zero: $x = 0$ or $x = K$. If the population size x is negative then $x' < 0$ and thus x is decreasing. If x is between 0 and K then $x' > 0$ and thus x is increasing. If x is larger than K, then $x' < 0$ and thus x is decreasing to carrying capacity K. Thus, $x = 0$ is an unstable equilibrium and $x = K$ is a stable equilibrium. This is illustrated with the arrows.

The carrying capacity can be interpreted as intraspecific competition of a plant species for resources, for example, nutrients or space. In the next example we introduce a grazer that feeds on these plants. The net rate of change of the plant population can be described by

$$\begin{pmatrix} \text{net plant population} \\ \text{growth rate} \end{pmatrix} = \begin{pmatrix} \text{gross plant population} \\ \text{growth rate} \end{pmatrix} - \begin{pmatrix} \text{grazing rate} \\ \text{by herbivores} \end{pmatrix}$$

We introduce a new parameter, the grazing intensity, c, which can be interpreted as the product of the number of grazers per unit area times the maximum amount of grass they each can graze per unit of time $\left(\frac{\sharp}{m^2} \times \frac{g}{d}\right)$. The actual grazing will also be dependent on the amount of grass. Only when the amount of grass per area is high the grazing rate will be limited by the herbivores, because it is reasonable to assume that they cannot eat more than needed: they are then saturated (this is called a Holling type 2 functional response). These assumptions can be represented very well by a rectangular hyperbola (or Michaelis Menten or Monod) equation. The combination of the logistic growth equation with the grazing term then yields differential Equation 2R.4

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = rx(t)\left(1 - \frac{x(t)}{K}\right) - c\frac{x(t)}{x(t) + h} \tag{2T.4}$$

where h is the amount of grass per unit area where the grazing rate is half as large as its maximum (half-saturation coefficient). Note that we assume that the number of grazers is not dynamical, i.e. their population is fixed and not dependent on the availability of food. The equilibria, where the net rate $\mathrm{d}x(t)/\mathrm{d}t = 0$, can be found by plotting the growth term (logistic growth rate) and the loss term (the grazing term) in one figure with the amount of grass per unit area on the x-axis, Figure 2T.6.
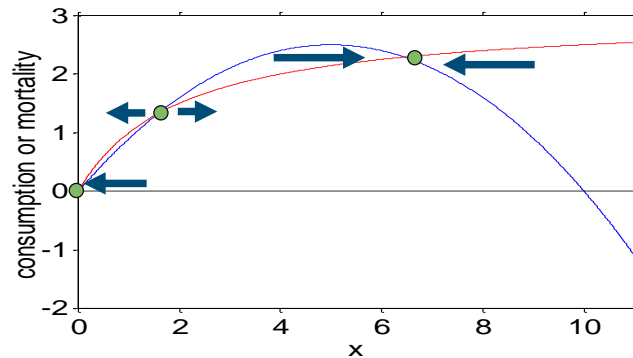

Figure 2T.6. The growth term (logistic growth rate), $rx(t)\left(1 - \frac{x(t)}{K}\right)$, and the loss term (hyperbolic consumption rate), $c\frac{x(t)}{x(t)+h}$, plotted against state $x(t)$. Their intersections are the equilibria.

The equilibrium points in this figure are the points where losses due to grazing exactly equal the growth of the plant population. Between the two rightmost equilibrium points, the gross plant population growth rate is higher than the grazing rate (parabola higher than Monod equation), which means that in this range the net growth of the plant population is positive. Before and after these equilibrium points, the consumption rate of the grass is higher than its growth rate, resulting in a negative net growth rate of the plant population (see arrows). The model of Equation 2T.4 is known as the Noy-Meir model, after the authors who introduced it in 1975 as a special case for predator (here a herbivore, for instance elephants) – prey (here grass) systems. Some other assumptions of the model are:

(i) the consumption rate of the herbivore population is the effective rate, because no distinction is made between, for example, juvenile and adult animals;

(ii) the growth rate of the vegetation (primary productivity) at any moment only depends on the quantity of vegetation present at that time; all other factors affecting growth are considered constant in time;

(iii) the increase in the low-biomass range in the logistic growth equation expresses the increase in photosynthetic capacity with increasing leaf area, and the levelling off to a plateau expresses self-interference effects such as shading and competition;

(iv) the half-saturation coefficient in the rectangular hyperbola determines the slope of the curve and can be interpreted as a grazing or searching efficiency of the herbivore, i.e. the extent to which it can find food to meet its requirements.

Analogous to the previous graphical analyses we can derive that the equilibrium point in the middle of Figure 2T.6 is unstable, whereas the other two equilibrium points are stable, as indicated by the arrows. The lower equilibrium is called the overgrazed state and the higher equilibrium the undergrazed state.

The fact that the grass population goes extinct below the unstable equilibrium point is comparable with the Allee effect described in Chapter 2P: however, the Allee effect applies to a single species that can go extinct below a so-called Allee extinction threshold, which holds for example for flamingos and penguins that do not mate below a certain animal density. In the grass population example there is a similar threshold, but now it is caused by the grazing function of another species.

An interesting question that can be asked about the Noy-Meir model is: what happens if grazing gradually increases? This implies that c in Equation 2T.4 becomes larger and larger due to more herbivores in the same area. Figure 2T.7-a shows a number of levels of grazing rates.
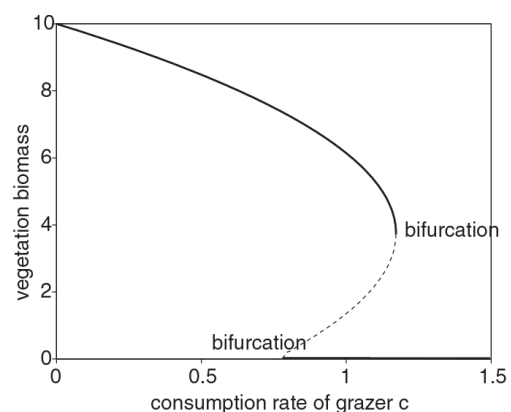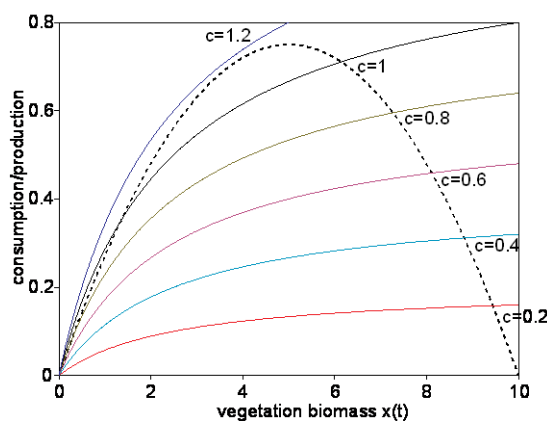
In Figure 2T.7-a the four lower hyperbolas ($c = 0.2; 0.4; 0.6; 0.8$) show levels of grazing with two intersection points for which the net rate of change in Equation 2T.4 equals zero, namely $x = 0$ and $0 < x < K$. The equilibrium at $x = 0$ is unstable, and the other one is stable. For the fifth grazing hyperbola ($c = 1$) there are three equilibrium values with the middle one unstable and the others stable, as before in Figure 2T.6. If grazing is increased further the two upper equilibrium values come closer to each other and eventually vanish ($c = 1.2$). We see that a gradual increase of the grazing rate, first shifts the amount of grass in the stable undergrazed equilibrium gradually to lower values, and if the hyperbola just surpasses the curve of the growth rate of the grass, implying overgrazing of the area, the grass goes extinct. So, if we plot the biomass in the vegetation at each of these stable equilibria against increasing grazing rate (Figure 2T.7-b), suddenly it will jump to zero. The point where this occurs is a bifurcation point, at $c = 1.2$, as at this parameter value the model behaviour changes significantly. Beyond this point there is only one stable equilibrium left. The interesting thing is that if the grazing pressure is diminished again, the grass population shifts back at a different level of grazing pressure (c). The population only returns to its old level if the grazing pressure is brought back to a much lower value, namely to the left bifurcation point at $c = 0.8$. More practically, it is better to abandon the grazing area to let the grass recover to a new critical mass before reintroducing herbivores, but in natural systems this is usually not possible. The jump at the right bifurcation point of the model, at $c = 1.2$, is really going to zero, which means that even if we lower the grazing pressure to very low values, the model grass can never recover in this deterministic model. In reality, however, there is always some grass left which cannot be grazed by the animals and the grazers will usually stop grazing if the biomass is below such a level. The model is thus too simple, but strongly helps to understand the dynamics of the system, in combination with some extra information about reality.

Figure 2T.7-b is called a bifurcation plot. It gives an overview of the effect of one parameter on the equilibria of the model. The parameter values where the equilibria change qualitatively (here one or two equilibria appear or vanish) are called bifurcation points.

References

Lotka, A. J. (1925). Elements of physical biology. Williams and Wilkins, Baltimore. [Reprinted in 1956: Elements of mathematical biology. Dover Publications, Inc., New York, New York].

Noy-Meir, I., 1975. Stability of grazing systems: an application of predator-prey graphs. J. of Ecology, Vol. 63, nr 2, pp.459-481.

Volterra, V., 1926. Fluctuations in the abundance of species, considered mathematically. Nature 118, pp.558-560.

Definitions used throughout the course

Allee effect: The fact that at low population densities the relative growth rate of a population is lower than at high population densities. This can for instance be due to a lack of opportunities for the individuals to find and mate with each other.

Alternative stable states: (see also: *multiple stable states*). Contrasting states to which a system may converge under the same external conditions. Which state is reached depends on the initial conditions of the state variables.

Attractor: A state or dynamic regime to which a system converges given sufficient simulation time. point or set of points in state space to which other points in this space contract if they are iterated through time.

Basin of attraction: The set of initial conditions from which you eventually reach the attractor.

Bifurcation point: Threshold in parameters (or conditions) that lead to a different quantitative bahavior of the model.

Carrying capacity: The value to which a population biomass or density grows in a certain environment. (i.e. the birth rates balance the death rates)

Catastrophic bifurcation: A bifurcation where an attractor disappears and the system is forced to move to another attractor. Such bifurcations are characterized by the fact that an infinitesimally small change in a control parameter can invoke a large change in the state of the system if it crosses the bifurcation. An example is the **fold bifurcation**.

Catastrophic transition: see critical transition.

Continuous time: Time in a model is not sampled or measured in steps, but essentially at every possible time the value of a state variable is known.

Critical transition: this kind of transition happens if a system shifts quite suddenly from one attractor to another

Determinism, deterministic: a world view in which everything is considered to be predictable as long as initial conditions and governing mechanisms are perfectly known.

Discrete time: Time in a model is sampled or measured in steps, for instance yearly or monthly steps.

Dynamical system: A description of the time evolution of a particular variable which is called a state variable

Equilibrium: A state is called an equilibrium when it holds that if you start exactly in that particular state you stay there indefinitely.

Extrinsic randomness: If a source of randomness is extrinsic, it lies outside of the system we are studying.

Fold bifurcation: A critical threshold in parameters where a catastrophic transition occurs. This bifurcation is part of a catastrophe fold.

Hysteresis: The tendency of a system to state in the same state if the conditions change. In the context of alternative stable states, the term refers to the fact that with increasing control parameters, the critical transition may occur at a different value than with a decreasing control parameter

Interspecific competition: Competition between individuals of different species.

Intraspecific competition: Competition between individuals of the same species.

Logistic growth: A simple model describing the growth of a population towards the carrying capacity.

Multiple stable states: see alternative stable states.

Non-linear dynamical system: a system of difference or differential equations that is not

linear, i.e. it includes non-linear interaction terms

Nullclines: **or** *zero-growth isoclines*: line where growth of a species is zero, so where equil. exists.

Parameter: A constant value that is used in a model. It usually doesn't change during a model run.

Perturbation: A sudden disturbance in the conditions of a system.

Phase plane analysis: The graphical analysis of nullclines in the plane of two state variables.

Repellor: An equilibrium point that is unstable, thus the opposite of an attractor.

Resilience: (Holling, 1973) The ability of a system to recover to the original state upon a disturbance, i.e. to the maximum perturbation it can take without causing a shift to an alternative stable state.

Saddle-node bifurcation: see fold bifurcation

Saddle point: A point in state space that is attracting in some directions and repelling in others.

Separatrix: The border between two basins of attraction.

Stable: Each equilibrium can be (locally) stable or unstable. An equilibrium is locally stable if when a system is originally in the equilibrium state and it is changed to a slightly different state, it would move back to the original equilibrium value.

Stability landscape: A representation of the stability of a model as a landscape of hills and valleys depending on parameter values.

State space: each point in this space completely specifies the state of the system

State variable: one of the variables that is considered to vary in time in a dynamical system, e.g. the number of rabbits per m^2 through time.

Trivial equilibrium: the real trivial equilibrium is where everything is zero. In some systems the trivial equilibrium is an equilibrium with one of the states equal to zero.

Transcitical bifurcation: A transcritical bifurcation is a parameter value where two equilibria cross. In biological systems it is often the point where one species goes extinct.

Transient dynamics: the trajectory of a system before it settles down on an attractor

Unstable: see stable (if it is not stable then it is unstable)

Zero-growth isoclines: see nullclines.