Chapter 8

Dynamical Systems

A recurring theme in this book is the use of mathematical models consisting of a set of differential equations to explore the behavior of physical systems as they evolve over time. Some examples we have encountered are the S-I-R epidemiological model, predator-prey systems, and the motion of a pendulum. We call such a set of differential equations a **dynamical system**. Dynamical systems play important roles in all branches of science. In this chapter we will develop some general tools for thinking about them, with particular emphasis on the kinds of geometric insight provided by the concepts of **state space** and **vector field**.

8.1 State Spaces and Vector Fields

The standard way to graph solutions

If you look back at the examples we've considered, many of them take the following form: we have two (or more) variable quantities x and y that are functions of time, and we want to find the nature of these functions. What we have to work with is a model for the way the functions $x(t)$ and $y(t)$ are changing—i.e., we are told how to calculate $x'(t)$ and $y'(t)$ whenever we know the values of x and y, and possibly t. From a given starting point, we typically used something like Euler's method to get values for x and y at times on either side of the starting value. We then graphed the solutions as functions of time—x against t and y against t.

In many instances, the rules determining $x'(t)$ and $y'(t)$ depend only on the current values of x and y, but not on the value of t, so that knowing the current state of the system (as specified by its x and y values) is sufficient to determine the future and past states of the system. Such systems are said to be **autonomous**. These are the only systems we will be considering in this chapter.

A new way to graph solutions: as trajectories in state space

In autonomous systems there is another way of visualizing the solutions that can be very powerful. Instead of plotting values of x and y as functions of time, we view these values as coordinates of a point in the x-y plane. As the system changes, the point (x, y) will trace out a curve in this plane. The point (x, y) is called a **state**, and the portion of the plane corresponding to physically possible states is called the **state space** of the system. The solution curves that get traced out in state space are called **trajectories**. By looking at three examples, we will see how this method of analysis can help us understand the overall behavior of a system.

There are a number of effective software packages available which can perform efficiently all the operations we will be considering, and one of them would probably be the most useful tool for exploring the ideas in this chapter. On the other hand, the basic numerical operations are quite simple, and it is easy to modify the programs developed earlier in the text to perform these operations as well. For those of you who enjoy programming, we will from time to time point out some of these modifications. It can be instructive to implement them in your own programs, and we urge you to do so.

Predator–Prey Models

In chapter 4.1, we looked at several models for the dynamics of a simple system consisting of foxes (F) and rabbits (R). Our first model was

$$\begin{aligned} R' &= .1R\left(1 - \frac{R}{10000}\right) - .005\,RF \quad \text{rabbits per month,} \\ F' &= .00004\,RF - .04\,F \quad \text{foxes per month.} \end{aligned}$$

When we started with the initial values $R(0) = 2000$ rabbits and $F(0) = 10$ foxes, Euler's method produced the following solutions for the first 250 months:

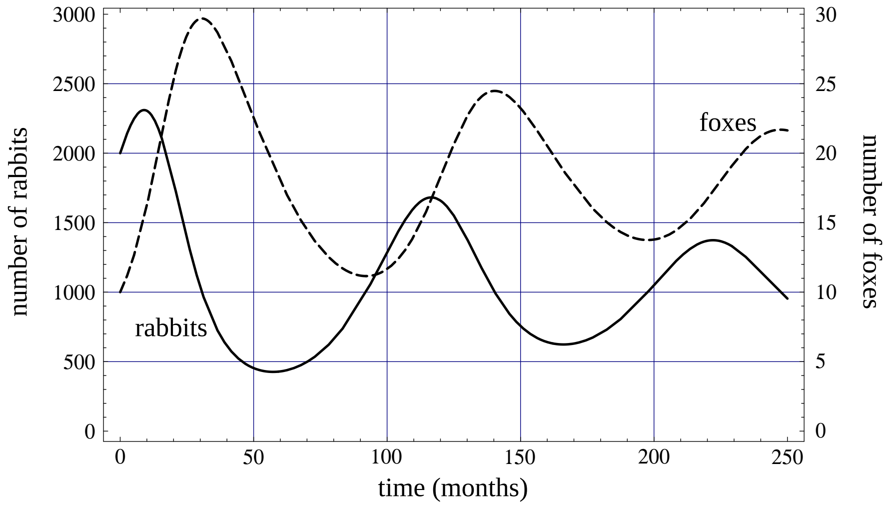

Two ways of representing the same solution graphically

Let's see how the same model looks when we express it in the language of state spaces. The state space consists of points in the R-F plane. For physical reasons our state space consists only of points (R, F) satisfying $R \geq 0$ and $F \geq 0$. That is, our state space is the first quadrant of the R-F plane together with the bounding portions of the R-axis and the F-axis.We can easily modify the program used to obtain the curves in the previous picture to plot the corresponding trajectory in the R-F plane. We only need change the specification of the dimensions of the viewing window and change the `plot` command to plot points with coordinates (R, F) instead of (t, R) and (t, F); all the rest of the calculations are unchanged. Here's what the same solution looks like when we do this:

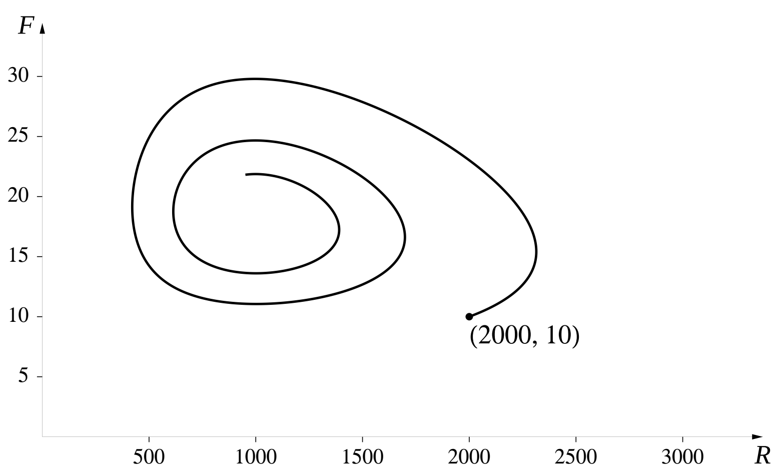

You should notice several things here:

- The trajectory looks like a spiral, moving in towards, but never reaching, some point at its center. We will see later (see page 469) how to determine the coordinates of this limit state.

- If we had started at any other initial state with $R > 0$ and $F > 0$, we would have gotten another spiral converging to the same limit (try it and see).

- From the trajectory alone, there is no way of determining the time at which the system passes through the different states. In part, this simply emphasizes that the succession of states the system moves through does not depend on the initial value of t, nor does it depend on the units in which t is measured—if t were measured in days or years, rather than in months, the trajectory would be unchanged.

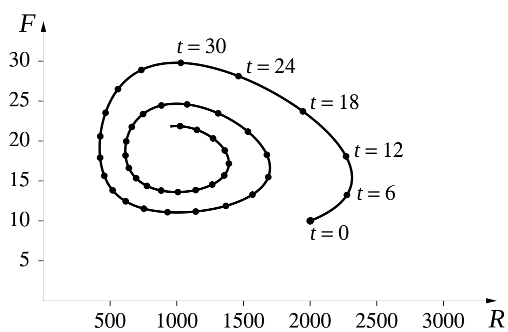

If we wanted to include some information about time, one way would be to label some points on the trajectory with the associated time value. If we label the points every 6 months, say, we would get the picture at the right. Note that the points are not uniformly spaced along the trajectory: the spacing is largest between points relatively far from the origin, where the values of R and F are largest. Moreover, the closer we come to the limit state, the tighter the spacing becomes.

The differential equations indicate how the state changes

Could we have foreseen some of this behavior by looking at the original differential equations? Since the differential equations give R' and F' as functions of R and F alone, for each point (R, F) in the state space we can calculate the associated values for R' and F'. Knowing these values, we can in turn tell in what direction and with what speed a trajectory would be moving as it passed through the point (R, F). Using our (by now) standard argument, in time Δt the change in R would be $\approx R'\Delta t$, while F would change by $\approx F'\Delta t$. We can convey this information graphically by choosing a number of points in the state space, and from each point (R, F) drawing an arrow to the point $(R + R'\Delta t, F + F'\Delta t)$. We would typically choose a value for Δt that keeps the arrows a reasonable size. Here's what we get in our current example when we choose a 16×16 grid of points in the region $0 \le R \le 3000$ and $0 \le F \le 30$, with $\Delta t = 1$.

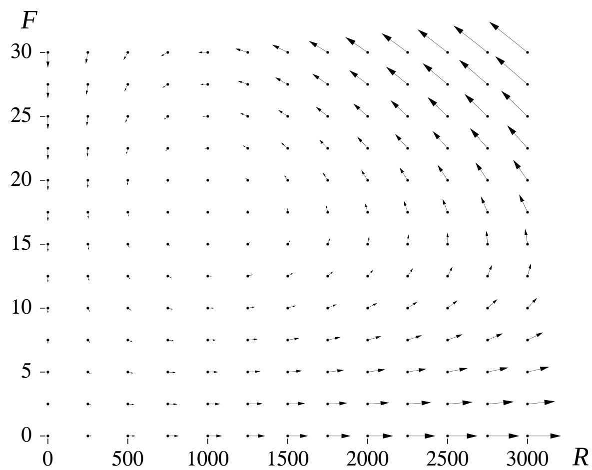

Arrows indicate the change in state...

Several things are immediately clear from this picture: the arrows suggest a general counter-clockwise flow in the plane; change is most rapid in the upper right corner; near the limit point of the flow and near the origin change is so slow that arrows don't even show up there.

...and trajectories are tangent to the arrows

Moreover, since the method used to construct the arrows is exactly the way Euler's method calculates the trajectories themselves, the solution trajectory through a given initial state is a curve in the state space which at every point is tangent to the arrow at that point. For instance, if we superpose the trajectory graphed on page 463 on the picture above, we get the following:

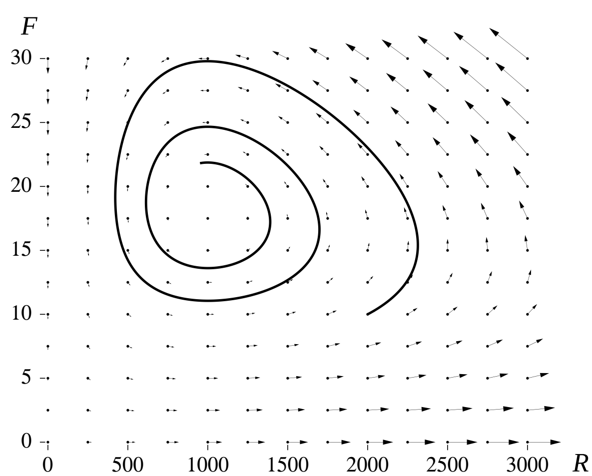

How to construct a geometrical visualization of a dynamical system

The net effect of this construction is thus to transform a problem in *analysis*—findinging a solution to a system of differential equations—into a problem in *geometry*—finding a curve which is tangent everywhere to a prescribed set of arrows. This correspondence between the analytical and the geometrical ways of formulating a problem is very powerful. Let's sum up the way this correspondence was established:

- We set up a **state space** for the system being studied. Each point—called a **state**—in the space corresponds to a possible pair of values the system could have.

- There is a rule which assigns to each point in the state space a **velocity vector**—which can be visualized as an arrow in the space based at the given point—specifying the rates at which the coordinates of the point are changing. The rule itself, which is just our original set of rate equations, is called a **vector field**. Geometrically, we can visualize the vector field as the state space with all the associated arrows.

- Solutions to the dynamical system correspond to **trajectories** in the state space. At every point on a trajectory the associated velocity vector specified by the vector field will be tangent to the trajectory. The existence and uniqueness principle for the solutions of differential equations—there is a unique solution for each set of initial values—is geometrically expressed by the property that every point in the state space lies on exactly one trajectory. The set of all possible trajectories is called the **phase portrait** of the system. For instance, part of the phase portrait of the system we have been considering appears below. We have drawn only a few trajectories—if we had drawn them all, we would have seen only a black rectangle since there is a trajectory through every point.

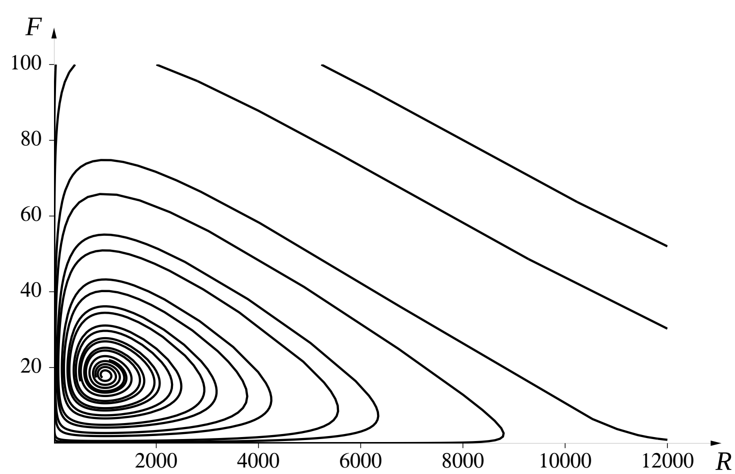

Simplifying the picture

There is almost too much detail in the picture of the vector field and the phase portrait. One way to see the underlying simplicity is to notice that the space is divided into four regions according to whether F' and R' are positive or negative. The *signs* of F' and R' in turn determine the *direction* of the associated velocity vector. For instance, if F' and R' are both positive, then F and R must both be increasing, which means the velocity vector will be pointing up and to the right, while if $F' > 0$ and $R' < 0$, the velocity vector will be pointing up ($F' > 0$) and to the left ($R' < 0$). Let's see which states correspond to which behaviors. Here are original rate equations:

$$R' = .1R\left(1 - \frac{R}{10000}\right) - .005\,RF \quad \text{rabbits per month,}$$
$$F' = .00004\,RF - .04\,F \quad \text{foxes per month.}$$

The equation for F' is slightly simpler, so we'll start there. We see that $F' = 0$ in exactly two cases:

1. when $F = 0$, or
2. when $.00004R - .04 = 0$, which is equivalent to saying $R = 1000$.

The first case simply says that if we are ever on the R-axis ($F = 0$), then we stay there—a trajectory starting on the R-axis must move horizontally. (If you start with no foxes, you will never have any at a later time.) The second case says that the value of F isn't changing whenever $R = 1000$. The set of points satisfying $R = 1000$ is just a vertical line in the state space. The condition that $F' = 0$ on this line can be expressed geometrically by saying that any trajectory crossing this line must do so horizontally (why?).

Divide the state space into regions

The remainder of the quadrant consists of two regions: one consists of all points (R, F) with $0 \le R < 1000$ and $F > 0$, the other consists of all points (R, F) with $R > 1000$ and $F > 0$. Moreover, since we've already accounted for all the points where $F' = 0$, it must be true that at every point of these two regions F' must be > 0 or < 0; F' can't equal 0 in either region. Further, within any one region F' must be always positive or always negative. If it were positive at some points and negative at others in a single region, there would have to be transition points where it took on the value 0, which we have just observed can't happen. (Be sure you see why this is so!) Thus to determine the sign of F' in an entire region, we only need to see what the sign is at one point in that region. For instance if we let $R = 2000$ and $F = 1$, we see that $F' = .08 - .04$, which is positive. Therefore we will have $F' > 0$ (fox population increasing) for any other state (R, F) with $R > 1000$. Similarly we can show that $F' < 0$ (fox population decreasing) if $0 \leq R < 1000$—the test point $R = 0$ and $F = 1$ is easy to evaluate. We could, of course, have arrived at the same conclusions through more formal algebraic arguments, which are fairly straightforward in this instance. In other problems, though, the "test point" approach may be the more convenient.

In exactly the same way, if we look at the first rate equation, we find that $R' = 0$ in two cases:

1. when $R = 0$, or
2. when $.1\,(1 - R/10000) - .005\,F = 0$. This is just the equation of a line, which can be rewritten as $F = 20 - .002\,R$.

The interpretations of these two cases are similar to the preceding analysis: any trajectory starting on the F-axis must stay on the F-axis; any trajectory crossing the line $F = 20 - .002\,R$ must cross it vertically, since $R' = 0$—the R-value isn't changing—there. Further, for any other state (R, F) we have $R' > 0$ if the point is below this line (the point $R = 1$ and $F = 0$ is a convenient test point where it's easy to see without doing any arithmetic that $R' > 0$), and $R' < 0$ if the point is above the line.

We can combine all this information into the following picture. We have drawn a number of velocity vectors along the lines where $R' = 0$ and $F' = 0$, with one or two others in each region.

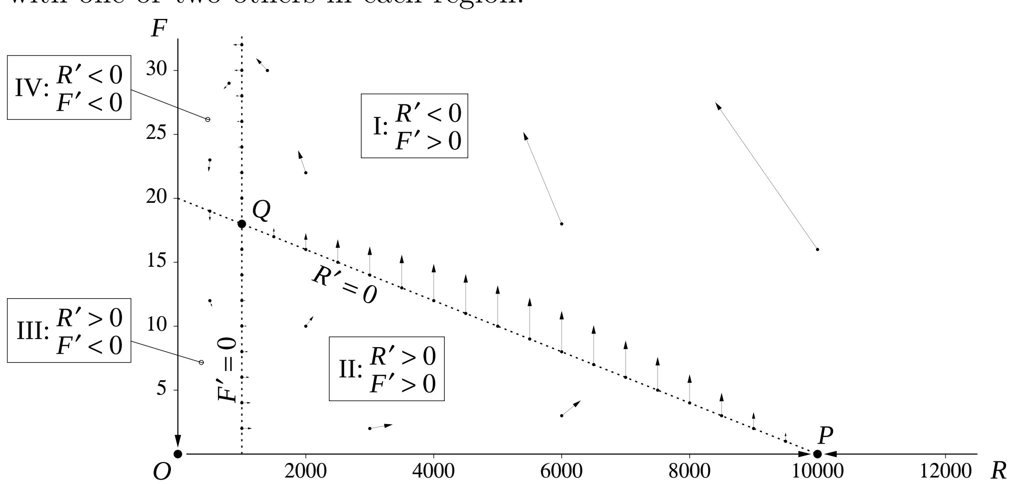

We see that the entire state space is divided into four regions:

1. Region I, above the line $F = 20 - .002\,R$ and to the right of the line $R = 1000$. Here $R' < 0$, and $F' > 0$, so all velocity vectors are pointing up and to the left.

2. Region II, below the line $F = 20 - .002\,R$ and to the right of the line $R = 1000$. Here $R' > 0$, and $F' > 0$, and all velocity vectors are pointing up and to the right.

3. Region III, below the line $F = 20 - .002\,R$ and to the left of the line $R = 1000$. Here $R' > 0$, and $F' < 0$, and all velocity vectors are pointing down and to the right.

4. Region IV, above the line $F = 20 - .002\,R$ and to the left of the line $R = 1000$. Here $R' < 0$, and $F' < 0$, and all velocity vectors are pointing down and to the left.

Notice that this diagram makes it clear what the limit state of the spirals is: it is the point $Q = (1000, 18)$ where the line $R = 1000$ and the line $F = 20 - .002\,R$ intersect. Notice that at Q both $R' = 0$ and $F' = 0$, so that if we are ever at Q, we never leave—the point Q is a trajectory all by itself. The points $O = (0, 0)$ and $P = (10\,000, 0)$ are the two other such point trajectories. While the typical trajectory looks like a spiral coming into the point Q, note that this picture contains three other "special" trajectories in addition to the point trajectories:

There are simple trajectories: three are just points...

...and three are straight line segments

- The F-axis for $F > 0$. The point $(0, 0)$ is *not* part of this trajectory.

- The portion of the R-axis with $0 < R < 10000$. Here the flow is toward the right, towards the point P.

- The portion of the R-axis with $10000 < R$. Flow is to the left, towards P, with movement being slower and slower as P is approached. Note that this is entirely separate from the preceding trajectory—you can't start at any point on one of them and get to any point on the other.

Equilibrium Points

Different kinds of equilibrium points

The three points O, P, and Q in the previous figure—single points which are also trajectories—are called **equilibrium points** for the system. If the ststem ever in such a state, it stays in it forever. Moreover, the system can't reach such a state from any other state (although it may be able to come very close). Nevertheless, the behavior of the system is not the same near the three points. If we zoom in on each of these points and draw some of the nearby trajectories, we get the following pictures:

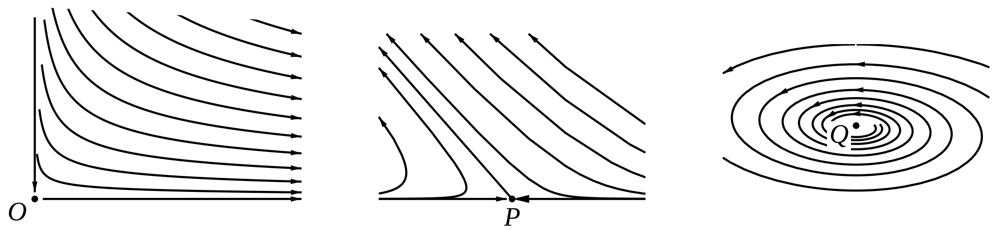

Points O and P look fairly similar—they would look even more alike if we crossed over into the negative R and negative F regions and included the trajectories there as well (impossible to do in the real world, but elementary in mathematics!). In both cases there is one direction from which trajectories come straight towards the point (in the case of O, this is the F-axis; for P this is the R-axis), and one direction in which trajectories move directly away from the point (the R-axis in the case of point O, and the line of slope $-.0092$ (we'll see how to find this later!) in the case of P). The remaining trajectories look sort of like hyperbolas asymptotic to these two lines. Equilibrium points of this sort are called **saddle points**. They are characterized by the property that there is exactly one direction along which the system can be displaced and still move back towards the equilibrium point. Displacements in any other direction get amplified, with the state eventually moving even further away.

Saddle point equilibrium

Point Q is quite different. If the state experiences a small displacement away from Q in any direction, over time it will move back towards Q. Such equilibrium points are called **attractors**, and Q is an example of a particular kind of attractor called a **spiral attractor**. In this example, Q is an attractor for almost the entire space—if we start with any point (R, F) with $R > 0$ and $F > 0$, the trajectory through (R, F) will eventually come arbitrarily close to Q and stay there. We will shortly see examples (see page 477, for instance) of attractors that draw from more limited portions of the state space.

Spiral equilibrium

Attractors and repellors

For future reference, we define here the concept of **repellor** and **spiral repellor**. Their vector fields look just like those for the attractors, but with all the arrows reversed. If the state experiences a small displacement from a repellor, over time this displacement will increase. We will see examples of a repellor in 8.4 It turns out that there is a relatively small number of kinds of equilibrium points that a system can have, and we will meet most of them in the next several examples. We will turn more systematically to the problem of identifying the kinds of equilibrium points in 8.2.

The Pendulum Revisited

In chapter 7 we analyzed the motion of a pendulum. Let's see how this analysis looks when translated into the language of state space. We first need to figure out what the appropriate coordinates are, which means deciding what information we need in order to specify the state of a pendulum. If you look back at the model in the last chapter, you will recall that the two variables we needed were the displacement x and the velocity v. Since x and v can potentially take on any values, our state space will be the entire x-v plane. As before, the dynamical system is specified by the equations

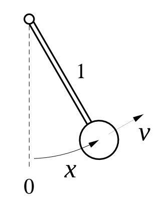

$$x' = v, \qquad v' = -\sin x.$$

Here is what the vector field for this system looks like:

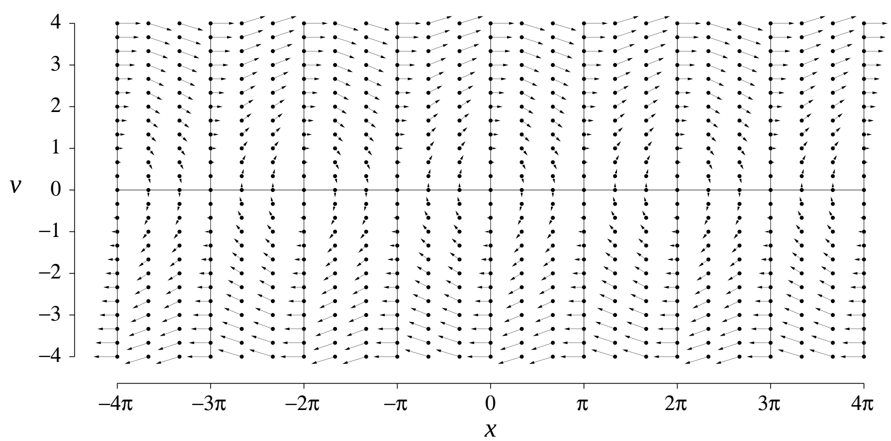

We have included in this diagram the lines where $v' = 0$ (the vertical lines at every multiple of π) and the line where $x' = 0$ (the horizontal line at $v = 0$). Note that the velocity vectors are horizontal on the lines corresponding to $v' = 0$ and are vertical on the line corresponding to $x' = 0$. The points where these two sets of lines intersect—all points of the form $(k\pi, 0)$ for k an integer—are the equilibrium points of the system. Let's sketch the phase

The equilibrium points

portrait of this system to see more clearly what's going on:

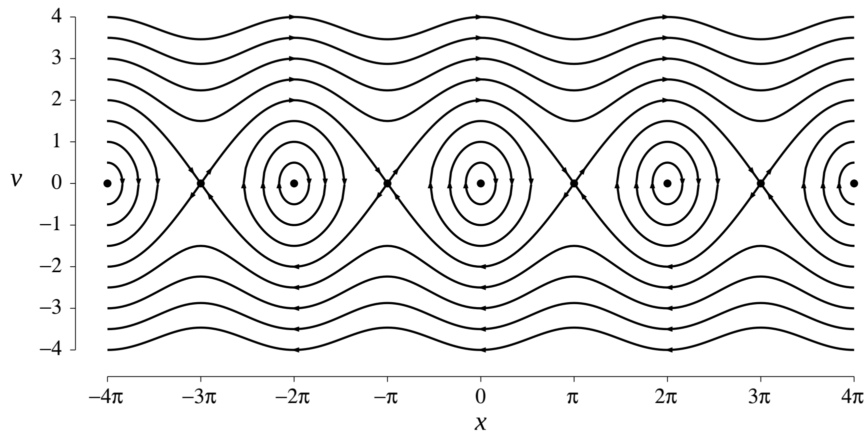

We see that there are several different kinds of trajectories:

Trajectories from left to right in the state plane correspond to $v > 0$

- There are the wavy trajectories moving from left to right across the top of the state space. Note that for these trajectories the value of v is always positive, and x just keeps increasing. These trajectories correspond to the cases where the velocity is great enough that the pendulum can go over the top, continuing to loop around counterclockwise (since x is increasing and x is measured in a counterclockwise direction) forever. Notice that v takes on its minimum value when x is an odd multiple of π, which is what we would expect, since the pendulum is at the top of its arc then. Similarly, v takes on its maximum value at the bottom of its arc—x an even multiple of π.

- The wavy trajectories moving from right to left across the bottom are similar, except that v is always negative. This corresponds to the pendulum wrapping around in a clockwise direction.

Oscillations of the pendulum correspond to closed loops in the state plane

- There are the closed loops. Here x oscillates back and forth between some maximum and minimum value symmetrically placed about an even multiple of π. These trajectories correspond to a pendulum swinging back and forth. The fact that some are centered at x-values other than 0 is due to the fact that the same position of the pendulum can be specified by an infinite number of values of x, all differing from each other by multiples of 2π.

A center: a neutral equilibrium

- There are the equilibrium points $(k\pi, 0)$, with k an even integer. This corresponds to the pendulum hanging straight down. If we perturb the system to a state slightly away from such a point, the pendulum swings back and forth, and the corresponding trajectory loops around the equilibrium point forever. The system neither comes back to the equilibrium point—the condition for an attractor—nor does it go wandering off even further away—the condition for a repellor. Such an equilibrium point is called a **center**. Notice that it is neither an attractor or a repellor; it is said to be a **neutral equilibrium**.

- There are the equilibrium points $(k\pi, 0)$, with k an odd integer, correspond to the pendulum balanced vertically. These are saddle points—if we perturb the system slightly with *exactly* the right v-value for the given x-value, the system will move back toward the vertical position; any other combination, though, will cause the pendulum to wrap around and around forever or to oscillate back and forth forever, depending on whether the v-value is greater than or less than the critical value.

- There are the trajectories connecting the saddle points. These correspond to cases where the pendulum has just enough velocity so that it keeps moving closer to the vertical position without either overshooting and wrapping around, or coming to a stop and reversing direction. In fact, these trajectories divide the state space: on one side of such a trajectory are points corresponding to states where the pendulum will wrap around, and on the other side are points corresponding to states where the pendulum will swing back and forth. Note that the saddle points are *not* part of these trajectories, and that each arc between saddle points is a separate trajectory—you can't get from a point on one of them to a point on another.

A connected curve in the phase portrait may be composed of more than one trajectory

First Integrals Again

In the case of the pendulum, we have another way of thinking about the trajectories. Recall that in chapter 7.3 we saw that, for any given initial conditions, the quantity $E = \frac{1}{2}v^2 + 1 - \cos x$ was constant over time. In the vocabulary of this chapter, if (x, v) is any state on the trajectory through (x_1, v_1), then it must be true that $\frac{1}{2}v^2 + 1 - \cos x = \frac{1}{2}v_1^2 + 1 - \cos x_1$. But this relation determines a curve in the x-v plane. We thus have an algebraic condition for each of the trajectories, which will depend on the initial values. In this example we could actually get an equation for each trajectory by first using the initial values to determine the energy $E = \frac{1}{2}v_1^2 + 1 - \cos x_1$.

First integrals can give equations of trajectories

We could then solve for v in terms of x by $v = \pm\sqrt{2(E-1+\cos x)}$ and plot the resulting function. (Whether we took the plus sign or the minus sign would depend on whether the pendulum was moving counterclockwise or clockwise.) Let's return to our previous sketch of the phase portrait and label some of the trajectories by their corresponding values of E:

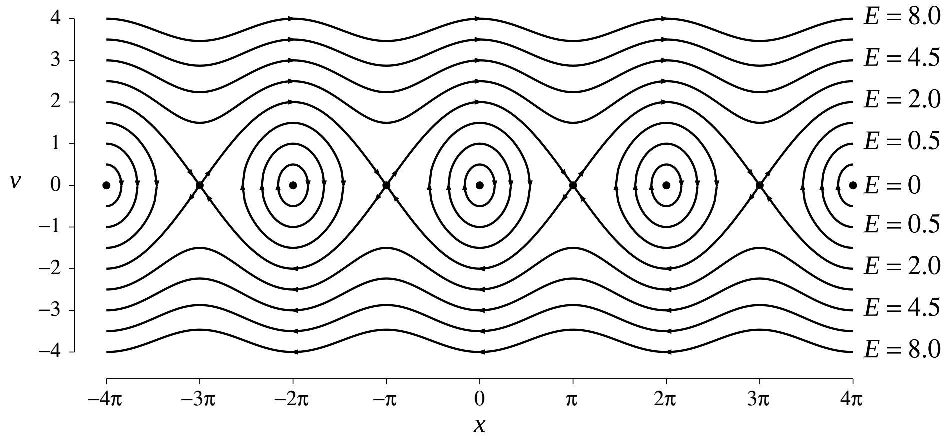

Note that for each value of $E \geq 0$ there is more than one trajectory having that value as its energy.

We can now characterize the different kinds of trajectories by their associated energy E:

- If $E > 2$ we get a trajectory extending from $x = \infty$ to $x = -\infty$ (or vice versa).
- If $0 < E < 2$, the trajectory is a closed loop.
- If $E = 0$, we get a neutral equilibrium point.
- If $E = 2$, we get either a saddle point equilibrium, or a trajectory connecting two such saddle points.

A Model for the Acquisition of Immunity

A mathematical model can be used to think about the feasibility of a proposed explanation

One of the roles of mathematical modelling is to allow researchers to explore possible mechanisms to explain an observed phenomenon. As an example of this, consider the phenomenon of immunity: for many infections, particularly those due to viruses, once you've been exposed to the disease your body continues to produce high levels of antibodies to the disease for the rest of your life, even in the absence of any further stimulation from the virus.

A capsule summary of the immune response: Vertebrates have a wide variety of specialized cells called *lymphocytes* circulating in their blood streams and lymphatic systems at all times. Each lymphocyte has the ability to recognize and bind with a specific kind of invading organism. The invader is called an *antigen*, and the neutralizing molecules produced by the responding lymphocytes are called *antibodies*. Prior to infection, the concentration level of a particular antibody is typically so low as to be undetectable, but the appearance of the antigen causes the system to respond by producing large quantities of the appropriate antibody. If the body can continue to produce high levels of antibodies, it will be immune to reinfection.

In their book *Infectious Diseases of Humans*, Roy Anderson and Robert May propose the following model as a possible mechanism for how antibody levels are sustained. Suppose that there are two kinds of lymphocytes (called *effector cells*) whose densities at time t are denoted by $E_1(t)$ and $E_2(t)$, with the type 2 cells being the potential antibodies for the disease in question. They assume further that new cells of type i ($i = 1$ or $i = 2$) are produced by the bone marrow at constant rates Λ_i and they die at a per capita rates of μ_i. They assume that each cell type is an antigen for the other—that is, contact with cell type 2 triggers cell type 1 to proliferate, and vice versa. They further assume that this proliferation response saturates to a maximum net rate which is dependent on the product of their respective densities. The following equations express this behavior:

$$dE_1/dt = \Lambda_1 - \mu_1 E_1 + a_1 E_1 E_2/(1 + b_1 E_1 E_2),$$
$$dE_2/dt = \Lambda_2 - \mu_2 E_2 + a_2 E_1 E_2/(1 + b_2 E_1 E_2).$$

Here the parameters Λ_i, μ_i, a_i, and b_i would have to be determined by experimental means. At this stage, though, when we are simply exploring to see if such a mechanism might account for the phenomenon of permanent immunity, we can try a range of values for the parameters to see how they affect the behavior of the model.

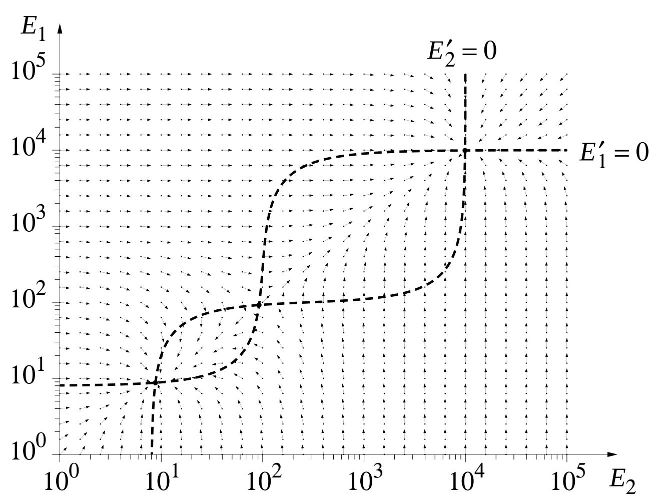

In the figure above, we have taken $\Lambda_1 = \Lambda_2 = 8000$, $\mu_1 = \mu_2 = 1000$, $a_1 = a_2 = 10$, and $b_1 = b_2 = 10^{-6}$.

There are a couple of features to notice about this graph:

Knowing the direction of the trajectories is often sufficient

1. Since the velocity vectors differ so much in their size, we have recorded only the direction of the velocity vectors, drawing all the arrows to be the same length. Thus we don't really show the vector field, but its close relative, the **direction field**. This is often a useful substitute.

Expressing graphical information over several different orders of magnitude

2. Since the range of values we want to represent is so great we have employed a common device from the sciences of plotting the values on a **log-log scale**. That is, we have plotted the values so that each interval spanning a power of 10—from 10^0 to 10^1, or from 10^3 to 10^4—gets the same space. This is equivalent to plotting the logarithms of the values on ordinary graph paper. This allows us to see effects that take place at different scales. If we hadn't done this, but had plotted this information on regular graph paper with the values running from 0 to 10^5, then some of our most interesting behavior—from 10^0 to 10^2—would be compressed into the lower left-hand corner of the graph, occupying only .001 of the vertical and horizontal scales.

We have included in the graph the two curves corresponding to all points satisfying $E_1' = 0$ and $E_2' = 0$ (note that these curves are *not* trajectories). These curves intersect at the three points $P_1 = (8.7689, 8.7689)$, $P_2 = (92.0869, 92.0869)$, and $P_3 = (9907.14, 9907.14)$, which are then the equilibrium points of this system. The points P_1 and P_3 appear to be attractors, while the point P_2 is a saddle point. In the next section (see page 487) we will see how to zoom in and look at the trajectories near each of these points to confirm this impression. Here is a picture of the phase portrait for this system.

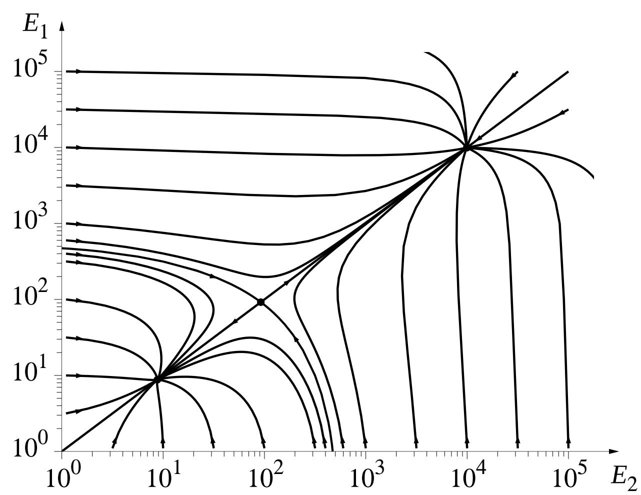

Note that neither P_1 nor P_3 is an attractor for the entire system. The **basin of attraction** for P_1 appears to be a region in the lower left of the graph, while the basin of attraction for P_3 is everything else. The boundary separating these two basins is formed by the two heavily shaded trajectories which come toward the point P_2 (since P_2 is a saddle point, there are only two such trajectories—every other trajectory eventually veers off and heads toward either P_1 or P_3).

We can now interpret this system in the following way. State P_1 represents the **virgin** or **resting state** of the system, with coordinate values on the order of magnitude of $E_i \approx \Lambda_i/\mu_i$, which would just be the steady state values we would have if there were no interactions between the two kinds of cells (why?). Note that after small perturbations (i.e., anything roughly less than a 10-fold increase of type 1 or type 2 cells) from P_1, the system will settle back to this resting state.

Now, though, suppose a viral pathogen appears which possesses an antigen which is identical to that expressed by cell type 1. This has an effect equivalent to moving vertically in the E_1-E_2 plane to a state which is now in the basin for P_3. As a result, the system immediately starts producing large quantities of type 2 cells (which are antibodies for the virus) very rapidly, the virus is wiped out, and the system settles into a new state—the **immune state**—P_3 and remains there. There are now so many type 2 cells permanently floating around the body that no further infection by the viral pathogen is possible. The only way the system can be switched back to state P_1 is if some other agent, such as radiation therapy or infection with an HIV virus, for instance, kills off large numbers of *both* the type 1 and type 2 cells, moving the system back into the basin of attraction for P_1. Just killing off large numbers of one type of cell won't move the system back to stateP_1—do you see why?

Exercises

Two-species interactions

We look at some variations of the predator-prey model. While the original context is given in terms of rabbits and foxes, similar models can be constructed for a variety of interactions between populations—not just predator and prey. The key features of the models are determined by the nature of the **feedback structure** between the populations. In the predator-prey models, the number of foxes has a negative effect on the growth rate of rabbits—the more foxes, the slower the rabbit population grows—while the number of rabbits has a positive effect on the growth rate of foxes. Can you think of other pairs of quantities whose interaction is of this sort? In the first problem we will look at several different models for predator-prey interactions. In the following three problems we will look at models for other kinds of feedback structures.

1. Below are four predator-prey models. In each model all the letters other than R and F are constant parameters. You can perform a general analysis, giving your answers in terms of the unspecified parameters a, b, c, etc., or, if you are more comfortable with specific values, perform the analysis using $a = .1$, $b = .005$, $c = .00004$, $d = g = .04$, $e = .001$, $f = .05$, $h = .004$, and $K = 10,000$. For each model you should carry out the following steps to sketch the vector field for the model in the first quadrant of the R-F plane. Compare your work with the steps that led up to the analysis of the vector field on page 468.

- Write down in words a justification for each rate equation—why is the model a reasonable one? What is it saying about the way rabbit and fox populations change?

- Draw (in red) the set of points where $R' = 0$, and mark the regions where $R' > 0$ and $R' < 0$.

- Draw (in green) the set of points where $F' = 0$, and mark the regions where $F' > 0$ and $F' < 0$.

- Mark the equilibrium points. What color are they?

- Sketch representative vectors of the vector field, and then sketch a couple of trajectories that follow these vectors. You might use a computer to verify your sketches.

- On the basis of your sketches make a conjecture about the stability of the equilibrium points.

a) The original Lotka–Volterra model, proposed independently in the mid-1920's by Lotka and Volterra. This model stimulated much of the subsequent development of mathematical population biology.

$$R' = aR - bR\,F,$$
$$F' = cR\,F - dF.$$

b) The Leslie–Gower model.

$$R' = aR - bRF,$$
$$F' = \left(e - f\frac{F}{R}\right) F.$$

c) Leslie–Gower with carrying capacity for rabbits.

$$R' = aR\left(1 - \frac{R}{K}\right) - bRF,$$
$$F' = \left(e - f\frac{F}{R}\right) F.$$

d) Another combination.

$$R' = aR\left(1 - \frac{R}{K}\right) - bRF,$$
$$F' = cRF + gF - hF^2.$$

2. **Symbiosis and mutualism**. Many flowers cannot pollinate themselves; instead insects like bees transport pollen from one flower to another. For their part, bees collect nectar from flowers and make honey to feed new bees. This sort of feedback structure in which the presence of each element has a positive effect on the growth rate of the other is called **symbiosis** or **mutualism** (there is a distinction made between these two interactions, but mathematically they are similar). Here is a model: B is the number of bees per acre, measured in hundreds of bees, while C is the weight of clover per acre, in thousands of pounds. Assume time to be measured in months.

$$B' = .1(1 - .01B + .005C)B,$$
$$C' = .03(1 + .04B - .1C)C.$$

a) Do these equations describe symbiosis? What terms account for symbiosis?

b) Each equation has a negative term in it. What aspect of reality is this term capturing?

c) Sketch the vector field for this system in the B-C plane. Find the equilibrium points, and mark them on your sketch.

d) Draw some trajectories on your sketch, and use them to determine the stability of the equilibrium points.

e) Suppose an acre of land has 10,000 pounds of clover on it, and a hive of 2,000 bees is introduced. (What are the values of $B(0)$ and $C(0)$ in this case?) What happens? Answer this question both by drawing a trajectory and by describing the situation in words.

f) Let a couple of years pass after the situation in part (e) has stabilized. Suppose the field is now mowed so only 2,000 pounds of clover remain on it. The bee–clover system is now at what point on the B-C plane? What happens now? Does the bee population drop? Does it stay down, or does it recover? Does the clover grow back?

g) This scenario is an alternative to part (f); it is also played out a couple of years after the situation in part (e) has stabilized. Suppose an insecticide applied to the clover field kills two-thirds of the bees. The insecticide is then washed away by rain, leaving the remaining bees unaffected. What happens?

3. **Competition** As a third kind of feedback structure, consider two species X and Y competing for the same food or territory. In this case each has a negative impact on the growth rate of the other. If we let x and y be the number of individuals of species X and Y, respectively, then the larger y is, the less rapidly x increases—and vice versa. Here is a specific model to consider:

$$x' = .15(1 - .005x - .010y)x,$$
$$y' = .03(1 - .004x - .005y)y.$$

The term $-.010y$ in the first equation shows explicitly how an increase in y reduces the growth rate x'. In the second equation $-.004x$ tells us how much X affects the growth of Y. Notice that Y affects X more strongly than X affects Y.

If x and y are both small, then the parenthetical terms are approximately equal to 1, so the equations reduce to

$$x' = .15x,$$
$$y' = .03y.$$

Thus, in these circumstances X's per capita growth rate is five times as large as Y's.

In the competition for resources will the growth rate advantage permit X to win the competition and drive out Y, or will the more adverse effect that Y has on the growth of X permit Y to win? Perhaps the two species will both survive and share the resources for which they compete. The purpose of this exercise is to decide these questions.

a) Suppose we start with $x = y = 10$. What are the two growth rates x' and y'? Is x' about five times as large as y' in this case? What are the approximate values of x and y after .5 time units have elapsed? Is X growing significantly more rapidly than Y?

b) How many equilibrium points does this system have, and where are they?

c) Sketch *and label* in the x-y plane the points where $x' = 0$ and where $y' = 0$. The vector field typically points in one of four directions: up and to the right; up and to the left; down and to the right; or, down and to the left. Indicate on your sketch the zones where these different directions occur and draw representative vectors in each zone.

[Note: Only three of the zones actually occur in the first quadrant; no vectors there point down and to the right.]

d) Sketch on the x-y plane the trajectory that starts at the point $(x, y) = (10, 10)$. Now answer the question: What happens to a population of 10 individuals each from species X and from species Y? In particular, does X gain an early lead? Does X keep its lead? Does either X or Y eventually vanish?

e) Is the outcome of part (d) typical, or is it not? Try several other starting points: $(x, y) = (150, 25)$, $(300, 10)$, $(200, 200)$, $(50, 200)$. Do these starting points lead to the same *eventual* outcome, or are there different outcomes? Use a computer to confirm your analysis.

f) Describe the type of each equilibrium point you found in part (a). Is any equilibrium an attractor?

4. **Fairer competition**. The vector field in question 3 shows that species X didn't have a chance: all trajectories in the first quadrant flow to the equilibrium at $(0, 200)$. We can attribute this to the strength of the adverse effect Y has on X—that is, to the size of the term $-.010y$ in the first equation when compared to the corresponding term $-.004x$ in the second equation.

Let's try to give X a better chance by increasing this term to $-.006x$. The equations become

$$x' = .15(1 - .005x - .010y)x,$$
$$y' = .03(1 - .006x - .005y)y.$$

a) Sketch and label in the x-y plane the points where $x' = 0$ and where $y' = 0$. Sketch representative vectors for the vector field. Mark all equilibrium points.

b) What happens to a population consisting of 10 individuals each from species X and species Y? Is the outcome significantly different from what it was in question 3? To get quantitatively precise results you will probably find a computer helpful.

c) What happens to a population consisting of 150 individuals from species X and 25 individuals from species Y? Is this outcome significantly different from what it was in question 3?

d) Is it possible for X and Y to coexist? What must x and y be? Is that coexistence *stable*; that is, if x and y are changed slightly, will the original values be restored?

e) Sometimes X wins the competition, sometimes Y. Mark in the x-y plane the dividing line between those starting points which lead to X winning and those which lead to Y winning.

f) Identify the type of each equilibrium point.

g) An often-articulated concept in ecology is the *principle of competitive exclusion*, which states that you can't have a stable situation in which two species compete for the same resource—one of them will eventually crowd out the other. Is the model you've been exploring in this problem consistent with such a principle?

5. **More on the Lotka–Volterra model**. The Lotka–Volterra model,

$$R' = aR - bRF,$$
$$F' = cRF - dF,$$

while it had a major impact on the development of mathematical biology, was found to be flawed in several important ways. The chief problem is that the equilibrium point $(d/c, a/b)$ is a neutral equilibrium point—given any starting state, the system would follow a closed trajectory. This in itself was all right and, in fact, stimulated a great many important investigations on whether or not cycles were an intrinsic feature of many populations. The difficulty was that there were so many possible closed trajectories—which one the system followed depended on where it started. A second difficulty, related to the first, is that there is a first integral for the Lotka–Volterra model. What is seen as a virtue in a physical system like the pendulum—since it is equivalent to the conservation of energy—is unrealistic in an ecological system, where there are almost certainly too many outside forces at work for any quantity to be conserved there. In the following exercises we will explore some of these behaviors. As before, you can either perform a general analysis of the model or use the specific parameter values $a = .1$, $b = .005$, $c = .00004$, and $d = .04$.

a) Sketch the vector field, together with some typical trajectories, in the rest of the R-F plane, including negative values. What happens to any trajectory starting at a state with a negative R or negative F value?

b) For this exercise you will need to go back to a computer program that implements Euler's method of approximating the trajectory by drawing a straight line segment from a point in the direction indicated by the velocity vector (commercial packages use fancier routines which accommodate for the kind of phenomena you are about to see!). Using the specific values for a, b, c, and d suggested above, starting from the point $(2000, 10)$ in the R-F plane, and using a time step $\Delta t = 1$, draw the first 500 segments of Euler's approximation to the trajectory. What does the trajectory look like? Would you think the trajectory was a closed loop on the basis of this result? How small does Δt have to be before the trajectory looks like it closes? Can you explain this phenomenon?

c) Using the same values for a, b, c, and d as in the preceding part, start at the point $(2000, 1)$ and use $\Delta t = 2$. This time calculate the first 1000 segments of Euler's approximation; what happens? (Your computer will probably give you some sort of overflow message.) Can you explain this? (Think about your answer to part (a).)

d) **Getting a first integral for the system** Show that the Lotka–Volterra equations imply that

$$\frac{R'}{R}(cR - d) = \frac{F'}{F}(a - bF).$$

Integrate this equation and show that the expression

$$cR + bF - d\ln R - a\ln F$$

must be a constant for all points on a given trajectory. If we know one point on the trajectory (such as the starting point), we can evaluate the constant.

e) Show that the function $f(R) = cR - d\ln R$ is decreasing for $0 < R < d/c$ and is increasing for $d/c < R < \infty$. Hence argue that for any given value of F there are at most two values of R giving the same value for the expression $cR + bF - d\ln R - a\ln F$. Hence conclude that the trajectories for the Lotka–Volterra equations can't be spirals, but must then be closed loops.

The pendulum

6. Suppose instead of an idealized frictionless pendulum, we wanted to model a pendulum that "ran down". One approach we might try is to throw in a term for air resistance. Let's see what happens when we add a term to the expression for v' which suggests that there is a drag effect which is proportional to the value of v—the larger v is, the greater will be the drag. Here are equations that do this:

$$\begin{aligned} x' &= v, \\ v' &= -\sin x - .1v. \end{aligned}$$

Perform a vector field analysis of this model, indicating the regions where the velocity vectors are pointing in the various combinations of up, down, right, and left. Try sketching in some trajectories. Where are the equilibrium points? What kinds are they?

The Anderson–May model

7. Consider $dE_1/dt = \Lambda_1 - \mu_1 E_1 + a_1 E_1 E_2/(1 + b_1 E_1 E_2)$. For what values of E_1 is it possible to find a value for E_2 making $dE_1/dt = 0$? Express your answer in terms of the parameters Λ_1, μ_1, a_1, and b_1. Is your answer consistent with the graph on page 475?

8. In the same book—*Infectious Diseases of Humans*—containing the previous model, Anderson and May propose another model to explain the acquisition of (apparently) permanent immunity. In this model there is just the

virus and the lymphocyte cells (effector cells) that kill the virus. We denote their populations at time t by $V(t)$ and $E(t)$. They propose the model

$$\begin{aligned} dE/dt &= \Lambda - \mu E + \varepsilon VE, \\ dV/dt &= rV - \sigma VE. \end{aligned}$$

Here Λ is the (constant) rate of background production of the lymphocytes by the bone marrow, μ is the per capita death rate of such cells, and r is the intrinsic growth rate of the virus if none of the specific lymphocytes was present. Both the increased production of the lymphocytes and the death of the virus are assumed to proceed at rates proportional to the number of their interactions, determined by their product.

a) Show that in the absence of any virus, the effector cells have a stable equilibrium of Λ/μ.

b) Perform a state space analysis of the vector field. Note that there will be two very different cases, depending on whether $\Lambda/\mu > r/\sigma$ or $\Lambda/\mu < r/\sigma$. In each case say what you can about the equilibrium points and the expected long-term behavior of the system.

c) Using parameter values $\Lambda = 1$, $\mu = r = .5$, and $\varepsilon = \sigma = .01$, and starting values $E = V = 1$ find the resulting trajectory. (The trajectory will be a spiral, but it moves in very slowly.)

d) How long, approximately, will it take the spiral to make one revolution? If this time, call it T, is roughly the same length as the lifetime of the infected individual, what will appear to be happening? It might help to plot both E and V as functions of time over the interval $[0, T]$.

8.2 Local Behavior of Dynamical Systems

A Microscopic View

One of the themes of this book has been the concept of the "microscope". When we zoom in on some part of a geometrical object, the structure typically becomes much simpler. In chapter 3 we used this approach to think about the behavior of functions. In this section we will use the same idea to analyze the behavior of a vector field and its phase portrait. There are two parts to this process:

Phase portraits under the microscope

1. We shift the origin of the coordinate system to center on the point we are interested in—we **localize**—and
2. We approximate both the vector field and its phase portrait by suitable linear approximations—we **linearize**.

To get a feel for how this works, let's go back and look at problem 4 on page 481 of the previous section. There we had two species X and Y competing for the same food source. We modeled the dynamics of this system by the equations

$$x' = .15(1 - .005x - .010y)x,$$
$$y' = .03(1 - .006x - .005y)y.$$

The phase portrait for this system looks like the following figure.

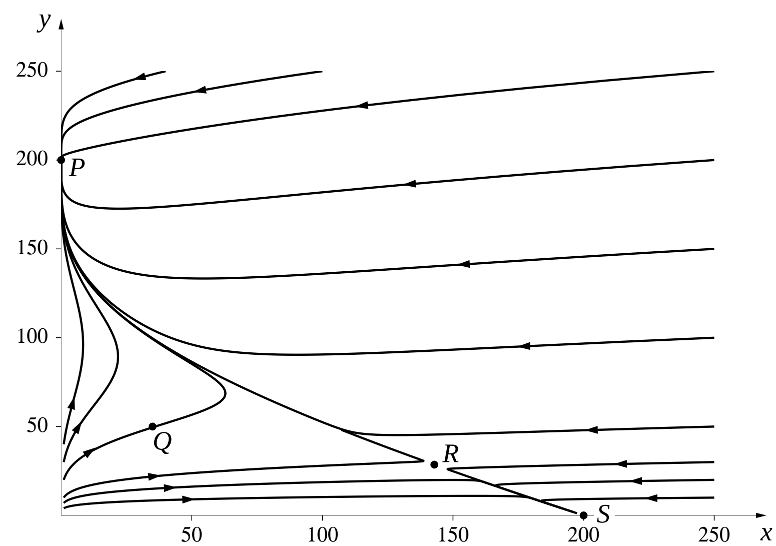

The three equilibrium points—$P = (0, 200)$, $R = (1000/7, 200/7)$, and $S = (0, 200)$—are indicated, together with a generic point $Q = (35, 50)$. Note that P and S are attractors and that R is a saddle point. As was the case with the Anderson–May model, there is a trajectory flowing away from R to each of the attractors. There are also two trajectories (not shown) flowing directly toward R and forming the boundary between the basins of attraction for P and S. We will see how to construct this boundary shortly (page 497).

Let's first zoom in on the point Q and see what the phase portrait looks like there. If we take the region ± 1 unit on either side of Q, we get the following phase portrait:

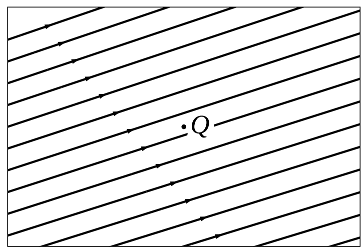

At this level, all the trajectories appear to be parallel straight lines. How could we have anticipated this picture? The first step in analyzing this phase portrait is to observe that since we are interested in its behavior near $Q = (35, 50)$, instead of working with the variables $x(t)$ and $y(t)$, we introduce new variables $r(t)$ and $s(t)$ which measure how far we are from Q:

$$\begin{aligned} r(t) &= x(t) - 35, \\ s(t) &= y(t) - 50. \end{aligned}$$

The effect of this transformation is simply to shift the origin to the point Q—the location of every point in the plane is now measured relative to Q rather than to the x-y origin. A point is close to the point Q if its r-s coordinates are small. Further, if we are given the r-s coordinates of a point, we can always recover the x-y coordinates, and vice versa—we can transform in either direction:

Shifting the origin

$$\begin{aligned} r &= x - 35, &\iff\quad x &= r + 35, \\ s &= y - 50, &\iff\quad y &= s + 50. \end{aligned}$$

Next, note that $r'(t) = x'(t)$ and $s'(t) = y'(t)$ so that the new variables change at the same rates as the old ones. We can now express our original differential equations in terms of the variables r and s by replacing x' by r', x by $r + 35$, y' by s', and y by $s + 50$. When we do this, we get

$$\begin{aligned} r' &= .15(1 - .005(r + 35) - .010(s + 50))(r + 35) \\ &= 1.70625 + .0225r - .0525s - .00075r^2 - .0015rs, \\ s' &= .03(1 - .006(r + 35) - .005(s + 50))(s + 50) \\ &= .81 - .009r + .0087s - .00018rs - .00015s^2. \end{aligned}$$

What we have accomplished by this is to transform a problem about trajectories near the point $(35, 50))$ in the x-y plane into a problem about trajectories near the origin in the r-s plane—we have **localized** the problem to the point we are interested in.

Near an ordinary point, a vector field is almost constant

The second step comes in analyzing the r-s system: since we are only interested in its behavior near the origin, we will be looking at values of r and s that are small. Under these circumstances, the contributions of the constant terms will far outweigh the contributions of any of the terms involving r and s. For instance, in our current example we are looking at a window that is ± 1 unit wide and ± 1 unit high around Q. In this window, the terms involving r or s are at most 3% of the constant term in the case of r', and a little over 1% in the case of s'. If we had used a smaller window, the contributions of the non-constant terms would be even less significant. This means that *near the r-s origin* the vector field for this system is well-approximated by the behavior of the related **constant linear system**:

$$r' = 1.70625,$$
$$s' = 0.81.$$

Note that 1.70625 and .81 are just the values of x' and y' at Q.

Near an ordinary point all trajectories look the same

In this linearized system, any change Δt in the time produces a change $\Delta r = 1.70625\Delta t$ in r, and a change $\Delta s = .81\Delta t$ in s. Thus the velocity vectors in the vector field near Q would all have the same length and would be pointing in the same direction, with slope $\Delta s/\Delta r = .81/1.70625 = .4747$. This in turn means that near Q all trajectories have the same slope and are traversed at the same speed.

We would see a similar picture—a family of parallel straight lines—whenever we zoom in on the phase portrait near any other ordinary (i.e., non-equilibrium) point (x_*, y_*) The vector field near such a point can always be approximated by a constant linear system of the form

$$r' = e,$$
$$s' = f,$$

where e and f are the values of x' and y' at (x_*, y_*). The trajectories of this approximating linear system will be lines of slope f/e.

Equilibrium points are different

Near an equilibrium point, the picture is more complicated. No matter how far in we zoom, the phase portrait never looks like a family of straight lines. For instance, here's what the picture looks like when we zoom in on $R = (1000/7, 200/7) \approx (142.857, 28.571)$:

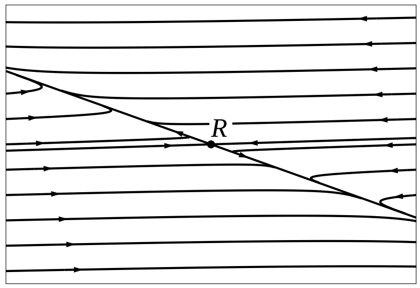

—if we zoomed in to a window 1/100-th the size of this one, the picture would be indistinguishable from this one.

Here we see four trajectories that look almost like straight lines—two coming directly towards R and two going directly away. All the other trajectories appear to be asymptotic to these two sets. On page 495 in the next section you will see how to find the equations of these asymptotes.

What happens when we linearize the vector field at R? As before, we first shift the origin so that it is centered at R by changing to coordinates r and s, where

$$\begin{aligned} r(t) &= x(t) - 1000/7, \\ s(t) &= y(t) - 200/7. \end{aligned}$$

When we then write the differential equations in terms of r and s, we get as before that $x' = r'$ and $y' = s'$ and

$$\begin{aligned} r' &= .15(1 - .005(r + 1000/7) - .010(s + 200/7))(r + 1000/7) \\ &= -.107143r - .214286s - .00075r^2 - .0015rs, \\ s' &= .03(1 - .006(r + 1000/7) - .005(s + 200/7))(s + 200/7) \\ &= -.00514286r - .00428571s - .00018rs - .00015s^2. \end{aligned}$$

Linear approximation of the vector field

This time, though, the constant term in the expression for both r' and s' is 0. This is because the point R was an equilibrium point, which meant that both x' and y', and hence r' and s', were 0 there. If we are considering only small values of r and s, though, say much smaller than 1, then the terms involving r^2 or s^2 or rs will be much smaller than the terms involving r and s alone. We can therefore simplify our equations at R by taking only the first powers of r and s, getting for the linearized system

$$r' = -.107143r - .214286s,$$
$$s' = -.00514286r - .00428571s.$$

In a similar fashion we could hope to explore the behavior of any other dynamical system about any of its equilibrium points by approximating the vector field there by a linear system of the form

$$r' = ar + bs,$$
$$s' = cr + ds,$$

for suitable constants a, b, c, and d.

We will see in section 8.3 how to use this linearized form of the vector field to discover many of the properties of equilibrium points.

How can we find values for the constants a, b, c, and d? If the differential equations specifying the rates of change of the variables are polynomials, then we can proceed as above:

- Shift the origin to the point we're interested in;
- Express the rate equations in terms of the new local variables;
- Throw away all the terms except the first degree terms.

This process requires some fairly tedious algebra. Moreover, what if the differential equations are not polynomials? Suppose, for instance, we wanted to study the local behavior of the Anderson–May model (page 475) at the saddle point $P_2 = (92.0869, 92.0869)$. Note that the differential equations are of the form

$$dE_1/dt = f_1(E_1, E_2),$$
$$dE_2/dt = f_2(E_1, E_2),$$

To linearize a vector field, linearize the functions that determine it

where f_1 and f_2 are the functions given in the text. But f_1 and f_2 are just functions, and we learned in chapter 3 how to construct locally linear approximations to them. This was, in fact, how we defined derivatives in the first place. Thus if E_1 changes by a small amount $\Delta E_1 = E_1 - 92.0869$, the function f_i will change by approximately $\partial f_i/\partial E_1 \times \Delta E_1$. Similarly, a small change $\Delta E_2 = E_2 - 92.0869$ will produce a change of approximately

$\partial f_i/\partial E_2 \times \Delta E_2$ in the function f_i. The total change in the function f_i can then be approximated by the sum of these changes:

$$\Delta f_1(E_1, E_2) \approx \frac{\partial f_1}{\partial E_1}\Delta E_1 + \frac{\partial f_1}{\partial E_2}\Delta E_2,$$
$$\Delta f_2(E_1, E_2) \approx \frac{\partial f_2}{\partial E_1}\Delta E_1 + \frac{\partial f_2}{\partial E_2}\Delta E_2.$$

General form for the local linearization at an equilibrium point...

But since P_2 is an equilibrium point, we have by definition that f_1 and f_2 are both zero there, so $\Delta f_i(E_1, E_2) = f_i(E_1, E_2) - f_i(P_2)$ is just $f_i(E_1, E_2)$. Further, if you look closely you will see that the quantity $\Delta E_1 = E_1 - 92.0869$ is identical with what we have been calling the local coordinate r, and $\Delta E_2 = E_2 - 92.0869$ is just the other local coordinate s. Thus, since $E_1' = r'$ and $E_2' = s'$, we have

$$r' = \frac{\partial f_1}{\partial E_1}\, r + \frac{\partial f_1}{\partial E_2}\, s,$$
$$s' = \frac{\partial f_2}{\partial E_1}\, r + \frac{\partial f_2}{\partial E_2}\, s,$$

where the partial derivatives are evaluated at P_2. Notice that there is nothing in this expression which is specific to this particular problem. The local linearization of any vector field at any equilibrium point will be in this form.

Finally, using the values given for the different parameters back on page 475, we can evaluate all the partial derivatives to get the specific local linearization for the point P_2:

$$r' = -94.5525\, r + 905.448\, s,$$
$$s' = 905.448\, r - 94.5525\, s,$$

We will see in the next section how knowing this form will allow us to find the boundary between the two basins of attraction.

...and at a generic point

For completeness, let's remind ourselves of what the local linearization would look like at a nonequilibrium point in the current formulation. The result is immediate and simple, using the analysis we used before. If Q is a generic point, then the local linearization consists of parallel lines, whose slopes are given by the constant rate equations

$$r' = f_1(Q),$$
$$s' = f_2(Q).$$

Exercises

1. Find the local linearizations at all the equilibrium points in exercises 2–4 at the end of the previous section.

2. a) The Lotka–Volterra equations

$$R' = aR - bRF,$$
$$F' = cRF - dF,$$

have an equilibrium point at $(R, F) = 00(d/c, a/b)$.

b) What is the local linearization there?

c) What is a striking feature of this linearization, and what is its physical significance?

d) The trajectories for the local linearizations turn out to be ellipses. If r and f are the local variables, find constants α and β such that the expression $\alpha\, r^2 + \beta\, f^2$ is constant on any trajectory.

3. Find the local linearization at the point P_1 in the Anderson–May model for the acquisition of immunity discussed in the previous section, using the parameter values given in the text on page 476.

4. Go back to the second Anderson–May model analyzed in problem 8 of the previous section (page 484). Using the parameter values given in part (c) there, find the local linearizations at all equilibrium points.

8.3 A Taxonomy of Equilibrium Points

An intuitive classification of equilibrium points

In the exercises and examples we have seen so far in this chapter, there have been several kinds of trajectories near equilibrium points: spirals towards and spirals away from the equilibrium, closed loops about the equilibrium, trajectories that looked vaguely like hyperbolas, and trajectories that seemed to arc more or less directly into or away from the equilibrium. It turns out that this rough classification covers virtually all the equilibrium behaviors we might encounter in a two-dimensional state space. There are many ways to demonstrate this, but we can accomplish almost everything with a couple of simple insights. We begin with a summary of the different kinds of equilibrium points, then turn to the question of devising ways to figure out from the equations what kind we are dealing with.

Suppose, then, that we are studying a two-dimensional dynamical system and that we have linearized the system at an equilibrium point. The point is either an attractor, repellor, saddle point, or neutral point. Attractors and repellors can be further subdivided according to whether they have one or two straight line trajectories, or whether their trajectories are spirals. Note that any attractor can be converted into a repellor simply by reversing the arrows, and vice versa (how do you accomplish this arrow reversal at the level of the defining differential equations?). If you reverse all the arrows at a saddle point, you get another saddle point. If you reverse the arrows at a neutral point, you get the same closed loops, but they are traversed in the opposite direction.

Here, then, is a listing of all the kinds of equilibrium points. There are five generic types. (*Generic* here means "general"; if you generate a random equilibrium point, it will almost certainly be one of these.) They are most easily categorized by whether or not they have **fixed line** trajectories—that is, trajectories which are straight lines going directly toward or directly away from the equilibrium point.

The existence or not of straight line trajectories and how to find them when they do exist is an instance of the so-called *eigenvector* problem. Analogous problems occur elsewhere in many parts of mathematics, physics, and even population biology. Being able to find such eigenvectors efficiently is an important problem in computational mathematics.

Nodes. Two pairs of fixed lines, all trajectories flowing toward the equilibrium (attractors) or away from it (repellors).

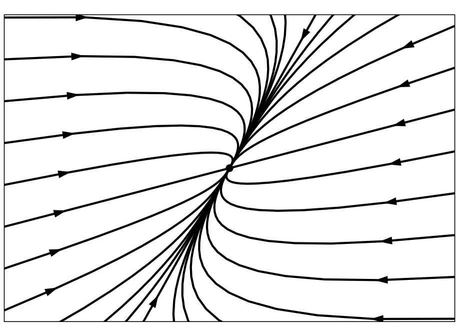

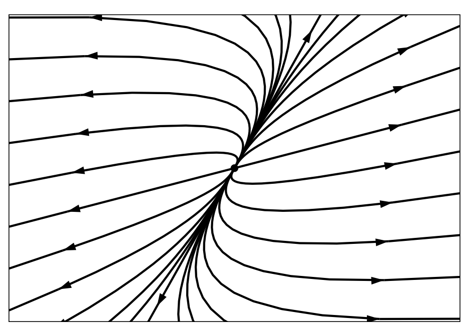

Spirals. No fixed lines, all trajectories spiraling toward the equilibrium (attractors) or away from it (repellors).

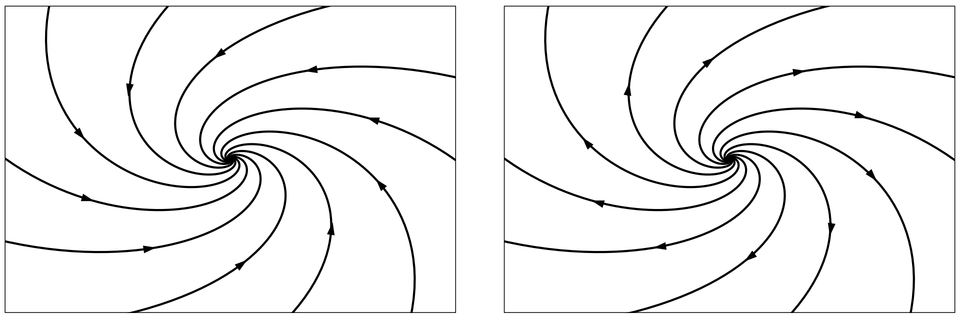

Saddle Point. Two pairs of fixed lines, with the flow along one pair being toward the equilibrium, and the flow along the other pair away from it. All other trajectories are asymptotic to these lines.

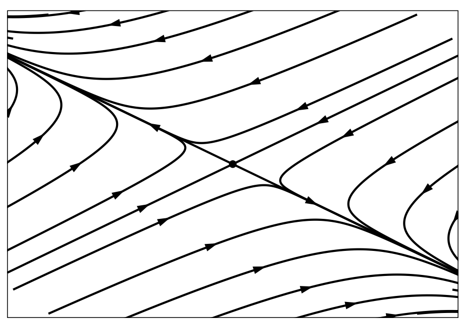

In addition to these five generic cases, there are three more types which arise under more specialized conditions:

Special Nodes. One pair of fixed lines, all trajectories flowing toward the equilibrium (attractors) or away from it (repellors).

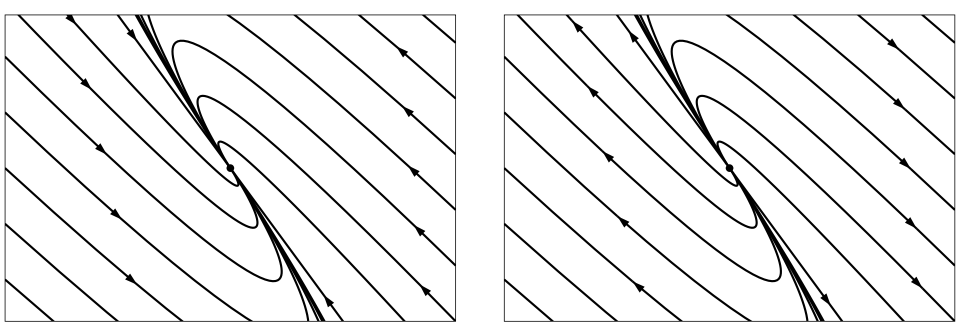

Center. No fixed lines, all trajectories flowing around the equilibrium in closed loops.

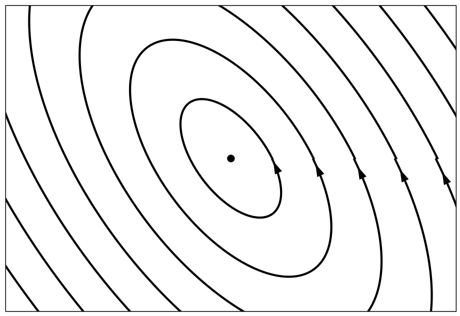

Except for a variety of highly specialized (or *degenerate*, in mathematical terminology) cases, examples of which are given in the exercises, the region near every equilibrium point will look like one of the above (although the exact shape may vary).

Clearly, it would be helpful to have an efficient way to determine whether or not fixed lines exist, and what their equations are if they do.

Straight-Line Trajectories

Given a dynamical system

$$r' = ar + bs,$$
$$s' = cr + ds,$$

how can we tell whether or not it has any straight-line trajectories? If $b = 0$, then the (vertical) line $r = 0$ is a trajectory. Otherwise, note that the line $s = mr$ will be a trajectory for this system provided the slope of the line—namely m—equals the slope of the vector field at every point (r, s) on the line. But the slope of the vector field at any point (r, s) is just s'/r', which in turn is equal to $(cr + ds)/(ar + bs)$. Since every point on the line of slope m is of the form (r, mr), what we are really asking, then, is whether there are any values of m which satisfy the equation

$$m = \frac{cr + dmr}{ar + bmr} = \frac{c + dm}{a + bm}$$

The condition for a fixed-line trajectory

To see how this works, let's return to the example of two competing species which we last looked at on page 486. There we zoomed in on the saddle point $R = (1000/7, 200/7)$ and found that the local linear approximation was

$$\begin{aligned} r' &= -.1071r - .2143s, \\ s' &= -.0051r - .0043s. \end{aligned}$$

If this system has a straight-line trajectory of slope m, then m must satisfy

$$m = \frac{-.0051 - .0043m}{-.1071 - .2143m},$$

which leads to the quadratic equation

$$.2143m^2 + .1028m - .0051 = 0,$$

which has roots

$$m = .0454 \quad \text{and} \quad m = -.5250.$$

Thus the lines $s = .0454r$ and $s = -.5250r$ are trajectories of the linear system. To be more exact, each of these lines is made up of three distinct trajectories: the portion of the line consisting of all points with $r > 0$, the portion with $r < 0$, and the origin (which is the saddle point R) by itself, which is always a trajectory in any linear system. To see whether flow along these trajectories is towards the origin or away from it, we could look to see where the lines lie in the state plane. It is just as simple, though, to try a test point. For instance, a typical point on the line $s = .0454r$ is $(1, .0454)$. When we substitute these values into the original rate equations, we find that

$$\begin{aligned} r' &= -.1071 \times 1 - .2143 \times .0454, \\ s' &= -.0051 \times 1 - .0043 \times .0454. \end{aligned}$$

We don't even need to do the arithmetic to be able to tell that both r' and s' are negative at this point, hence both r and s are decreasing, which means that on the line of slope .0454 movement is towards the origin. Similarly, on the line of slope $-.5250$ the flow is away from the origin. Finally, it turns out (as is the case with every linear system with straight-line trajectories) that every other trajectory is asymptotic to these lines.

The crux of this approach was the use of the quadratic formula. Of course, it may happen—and we will see examples in the exercises—that when we try the same approach on another system we find there are no real roots to the equation. This means that there are no fixed lines, so that trajectories must be spirals or closed loops.

Attractors and Basins of Attraction

One byproduct of the analysis in the previous section is that it gives us a technique for sketching the boundary separating two basins of attraction. Let's continue with the previous example to illustrate how this is done. We observed that the boundary between the two basins was formed by the two trajectories coming directly into the saddle point R between the two attractors P and S. We have just seen that near R these two trajectories looked like the straight line of slope .0454. We can therefore take a point on this line on each side of R and run the system backward (if we go forward, we simply approach R) in time to reconstruct the trajectories, and hence get the boundary of the basins of attraction.

Exercises

1. In this exercise we look at a number of different linear systems to see what kinds of trajectories we get. In each case you should sketch the trajectories. Do this as before by first identifying the regions in the plane where $r' = 0, r' > 0$, and $r' < 0$, and similarly for s'. Then sketch trajectories consistent with this information. You might want to use a graphing program to check any answer you're unsure of.

a) $r' = 4r + s, \qquad s' = 2r + 3s.$

b) $r' = 4r + s, \qquad s' = -2r + 3s.$

c) $r' = 2r + 3s, \qquad s' = 4r + s.$

d) $r' = -4r + 4s, \qquad s' = 2r + s.$

e) $r' = -.4r - 4s, \qquad s' = 2r - .5s.$

f) $r' = -.4r - 4s, \qquad s' = 2r + .4s.$

g) Make up and analyze four more linear systems.

2. If you start with a given linear system and consider the related system in which all the coefficients are four times as big, how do the trajectories change?

3. If you start with a given linear system and consider the related system in which all the coefficients have their signs reversed, how do the trajectories change?

4. a) Use the quadratic formula to find the general solution to the equation

$$m = \frac{c + dm}{a + bm}.$$

b) In exercises 1, 3, and 4 in the previous section you found local linearizations at the equilibrium points of a number of examples discussed earlier. Determine which of these have straight-line trajectories and which do not. For those that do, find the equations of the lines and determine for each line whether the flow is towards the origin or away from it.

c) What is the general condition for a linear dynamical system to have straight-line trajectories?

5. Make up a system that has the lines of slope ± 1 as trajectories.

6. What is the condition for a system to have exactly one fixed line? Construct a couple of systems that have only one fixed line and sketch their phase portraits.

7. **Degeneracy** The analysis developed in this section implicitly assumed that in the local linearization, at least one of the coefficients in each of the expressions for r' and s' was non-zero. If this is not true, then many more possibilities open up. The following two systems have the origin as their only equilibrium point. In each case, write down the local linearization and draw in the trajectory pattern for the linearized system. Notice that the linearized systems have more than one equilibrium point. Then do the standard phase plane analysis for the original system—identify the regions in the plane where $r' = 0$and where $s' = 0$, and specify what the direction field is doing in the rest of the plane, as usual. Sketch in some typical trajectories. Comment on the connections between the linearized and unlinearized forms.

a) $r' = r^2, \qquad s' = -s$. You should see sort of a hybrid between a saddle point and an attractor here.

b) $r' = r^2 + s^2, \qquad s' = r$.

8. a) Use the technique presented at the end of this section (page 497 to graph the boundary between the two basins of attraction.

b) In the same way, construct the boundary between the two basins in the competing species model we've been discussing—problem 4 on page 481.

Distance from the Origin

Another way to distinguish between different kinds of trajectories is to see how their distance from the origin varies over time. For saddle points the distance will first decrease and then increase. For spiral attractors and nodal attractors the distance may be always decreasing, or it may fluctuate, depending on how flat the trajectory is.

Again, let's look at a general linear system

$$\begin{aligned} r' &= ar + bs, \\ s' &= cr + ds. \end{aligned}$$

onsider the system moving along some trajectory in r-s space. At time t it will be at a point $(r(t), s(t)$, situated at a distance $d(t) = \sqrt{r(t)^2 + s(t)^2}$. We would like to know how the function $d(t)$ behaves. Is it always increasing? Always decreasing? Or does it have local maxima and minima? To answer this we need to know if $d'(t)$ is ever $= 0$, or if it is always positive or always negative. We can simplify our calculations if we look at the square of the distance: $D(t) = d(t)^2 = r(t)^2 + s(t)^2$. The function D will be increasing and decreasing at exactly the same points as the function d, and it's easier to work with.

9. a) Show that

$$\begin{aligned} D'(t) &= 2r(t)r'(t) + 2s(t)s'(t) \\ &= 2[r(ar + bs) + s(cr + ds)] \\ &= 2[ar^2 + (b + c)rs + ds^2]. \end{aligned}$$

b) Show that if we look at points on the line of slope m, so that $s = mr$, we will have $D'(t) = 0$ there if and only if

$$a + (b + c)m + dm^2 = 0\,.$$

c) Use the quadratic formula to conclude that this happens precisely where

$$m = \frac{-(b + c) \pm \sqrt{(b + c)^2 - 4ad}}{2d}\,.$$

d) Show in particular, if $(b+c)^2 - 4ad < 0$, there are no solutions to $D'(t) = 0$, and the distance must always be strictly increasing along all trajectories, or strictly decreasing along all trajectories.

e) Return to the example

$$r' = -.1071r - .2143s,$$
$$s' = -.0051r - .0043s,$$

and find the equations of the two lines where trajectories pass closest to the origin. These lines will not be trajectories themselves. Their significance is that the 'vertices' of all the trajectories will lie along them.

10. Choose four of the exercise in the first part of this section and analyze them to see where (and whether) trajectories have a closest approach to the origin.

11. a) Use the results of this section to construct a dynamical system whose trajectories are spirals that are always moving away from the origin.
b) Use the results to construct a dynamical system whose trajectories are flattened spirals, so that the distance from the origin, while increasing overall, has local maxima and minima.

12. It turns out that trajectories which form closed loops should really be considered as a special kind of spiral. In fact, a flattened spiral will close up precisely when the two directions in which the distance is a maximum or minimum are perpendicular to each other. Express this as a condition on the coefficients a, b, c, and d in the dynamical system.

13. Write down the equations of some dynamical systems that will have closed orbits.

8.4 Limit Cycles

With this analysis of the behavior of vector fields near equilibrium points, we now know most of the possibilities for the long-term behavior of trajectories. The one important phenomenon we haven't discussed is **limit cycles**. To see an example of this, let's return to May's predator–prey model we first encountered in chapter 4. If $x(t)$ and $y(t)$ are the prey and predator populations, respectively, at time t, then the general form of May's model is

$$x' = ax\left(1 - \frac{x}{b}\right) - \frac{cxy}{x+d},$$
$$y' = ey\left(1 - \frac{y}{fx}\right);$$

the parameters a, b, c, d, e and f are all positive.

Using parameter values of $a = .6$, $b = 10$, $c = .5$, $d = 1$, $e = .1$ and $f = 2$, let's take several different starting values and sketch the resulting trajectories. Here's what we find:

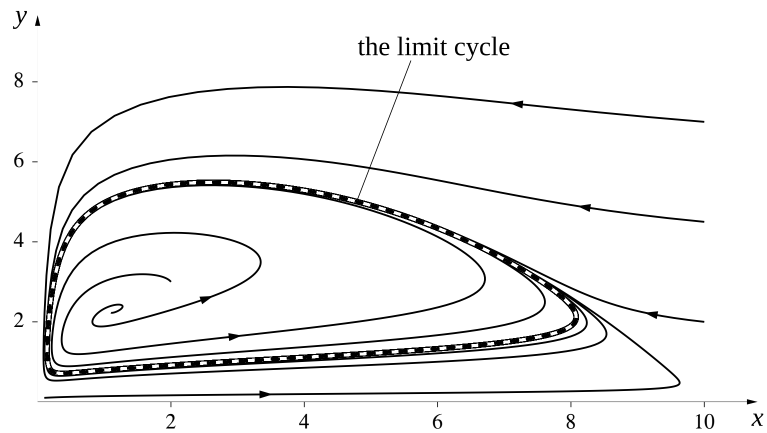

Notice that no matter where we start, the trajectory is apparently always drawn to the closed loop shown in dashes above. This loop is an example of an **attracting limit cycle**. As usual, we could reverse all the arrows in our vector field, in which case this example would be converted to a **repelling limit cycle**.

A limit cycle is very different from the kind of behavior we saw in the neighborhood of a neutral equilibrium point called a center. Around a center there is a closed loop trajectory through every point: displace the state slightly, and it would move happily along the new loop. If the state is on an attracting limit cycle, though, and you displace it, it will move back toward the cycle it started from. For this reason limit cycles make very good models for cyclic behavior, whether it is in the firing of neurons or population cycles of mammals.

Limit cycles give models for cyclic behavior

The size of the limit cycle, and even its very existence, depends on the specific values of the parameters in the model. If you change the parameters, you change the limit cycle. If you change the parameters enough, the limit cycle may disappear all together. (See the exercises.)

A result proved early in this century is the **Poincaré–Bendixson Theorem** which says that equilibrium points and limit cycles are as complicated as dynamical systems in two variables can get. Once we pass to three variables, the situation becomes much more complicated. Many of the phenomena associated with such systems have been discovered only within the past 50 years, and their exploration is a subject of continuing research. In the next section we will give a brief introduction of some of the new behaviors that can arise.

Exercises

1. Using the parameter values given in the text, find the coordinates of the equilibrium point at the center of the limit cycle and show that it is, in fact, a repellor.

2. May's model is interesting in that it exhibits a phenomenon known as **Hopf bifurcation**. Namely, the existence of a limit cycle depends on the values of the parameters. Choose one of the parameters in May's model and try a range of values both larger and smaller than in the example we've worked out. At what value does the limit cycle disappear? When this happens, the equilibrium point inside the cycle has become an attractor. Can you work out analytically when this happens?

8.5 Beyond the Plane: Three-Dimensional Systems

Up to now we have worked with dynamical systems in which there are only two interacting quantities. We have thought of the two quantities as specifying a point in the state space, which we think of as some subset of the plane. The dynamical system defined a vector field on the state space. These geometric notions carry over to dynamical systems involving more than two interacting quantities.

In particular, if we have a dynamical system consisting of three interacting quantities, then we think of the values of the three quantities as specifying a point (or state) in three-dimensional space. So, for instance, if we have an ecological system consisting of three species, then we think of the numbers x, y, z of each of the three species as specifying a point (x, y, z) in space. The set of all possible points or "states" is the set of points

$$\{(x, y, z) : x \geq 0, y \geq 0, z \geq 0\}$$

that constitute the "first octant" in Cartesian 3-space. We think of the dynamical system as a vector field: that is, as a rule which assigns to each point of the state space a vector. As in the case of the plane, we can define equilibrium points, trajectories, limit cycles, attractors, and the like.

In three-space there is a much wider range of behavior possible, even in the case of equilibrium points. We do, of course, have point attractors and repellors: all trajectories near a point attractor flow towards the attractor and all trajectories near a point repellor flow away from the repellor. However, a greater range of combinations is possible: an equilibrium point can attract all points along some plane, but repel all other points. Or the equilibrium point could be a center, surrounded by closed orbits lying in some plane which attract trajectories off the plane.

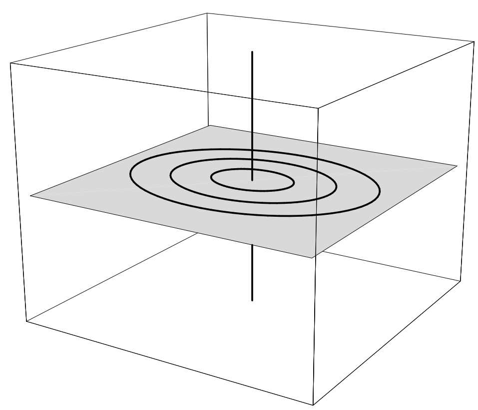

It is worth pointing out that we can represent the two-dimensional systems we've been exploring so far in this chapter in three dimensions by introducing a time axis. This has the effect of 'unwinding" the trajectories by stretching them out in the t-direction: closed trajectories become endless coils, equilibrium points become straight lines parallel to the t-axis, and so on.

The analytic tools we introduced to find and explore the nature of equilibrium points in two-dimensional systems carry over to three dimensions. In particular, in investigating the nature of an equilibrium point analytically, we first localize the system at the equilibrium point and linearize. The behavior of the linearized system can then be explored using analogues of the techniques introduced in the previous section (or using simple linear algebra).

There are also, of course, limit cycles, which can be attractors, repellors of a mix of the two (attracting, for example, all trajectories on a plane, but repelling all trajectories off the plane) in three-dimensional systems. As in the case of dynamical systems in the plane, attracting limit cycles in a three-dimensional system signal stable periodic behavior.

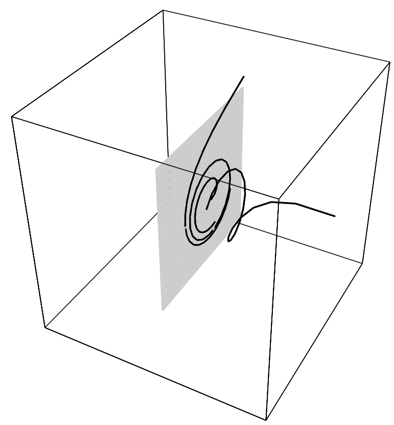

However, more complicated types of periodic behavior are possible in the three-dimensional case: we could, for example, have an attracting torus in the state space.

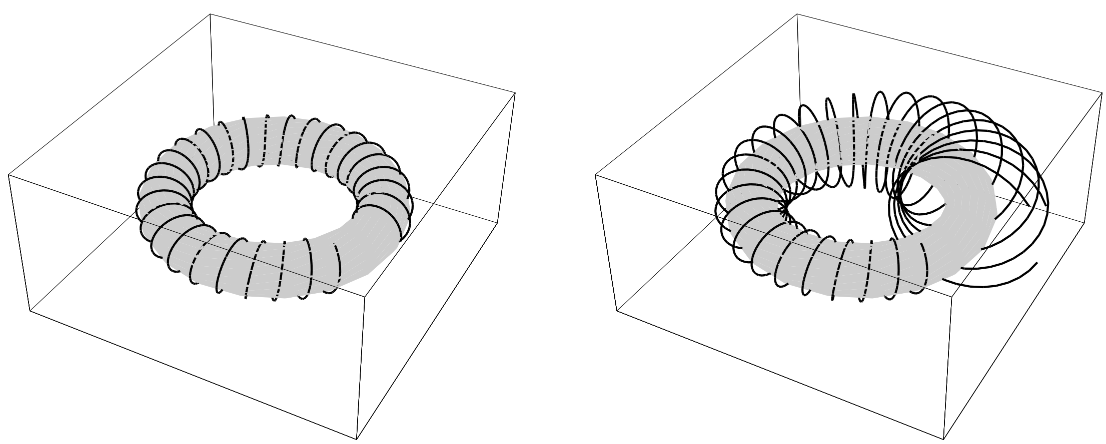

In this case, the behavior of the states does not settle down to periodic behavior, but a behavior which is approximately periodic (often called **quasi-periodic**). In the plane, there is a well studied phenomenon called Hopf bifurcation in which changing the parameters in a dynamical system can cause an attracting fixed point to become a repellor surrounded by a stable attracting limit cycle. Such dynamical systems arise in modelling situations in which a state begins to oscillate. In three dimensions we also sees the same sort of phenomenon in which an attractor can give birth to an attracting limit cycle. However, there are also three-dimensional systems in which varying the parameters results in an attracting limit cycle becoming a repelling limit cycle enclosed by an attracting torus (this is also called Hopf bifurcation and is frequently encountered in applications).

These sorts of behavior are relatively straightforward generalizations of behavior in the plane. At the turn of the century, Poincaré realized that simple three-dimensional systems could have exceedingly complicated trajectories which exhibit behavior totally unlike any two-dimensional trajectory. Discoveries in the last three decades have made it clear that qualitatively new types of *attractors* (not just trajectories) can exist in three-dimensional systems with even very simple equations. The most famous such attractor was discovered by a meteorologist, Edward Lorenz, in the course of using dynamical systems to model weather patterns. He discovered a class of simple systems with an attractor which corresponded to behavior which was in no sense periodic. An example of such a system is

$$\begin{aligned} x' &= -3x - 3y, \\ y' &= -xz + 30x - y, \\ z' &= xy - z. \end{aligned}$$

All trajectories of the system entered a bounded region of the state space and tended towards a clearly defined geometrical object (resembling a butterfly). But along the attractor, nearby points followed trajectories which rapidly diverged from one another. Below, we have sketched two views of a trajectory beginning at (0,1,0) of the system above.

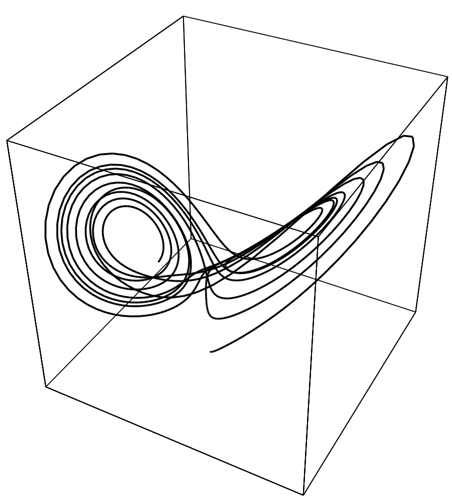

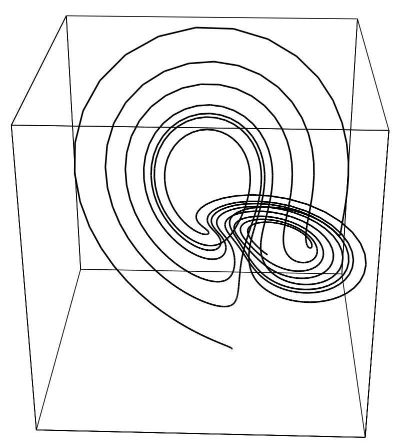

As Lorenz noted in the paper describing his discovery (Deterministic Nonperiodic Flow, *J. Atmos. Sci.*, **20** 130 (1963)), this divergence of trajectories along an attractor has astonishing practical implications. It means that that the trajectories of nearby points in state space could (and would) wind up following very different paths along the attractor. Since we never know initial conditions exactly (and even if we did, a computer truncates decimal expansions of the coordinates of any point, effectively replacing the point with a nearby point), this means that long-term predictions using a model possessing such an attractor are impossible. In other words, although the future is completely determined by a dynamical system given an initial state, it is unknowable in systems of the sort discovered by Lorenz, because initial states are never known exactly in practice. Such systems are called **chaotic** and attractors which are not points, limit cycles or tori are called **strange attractors**. These systems have been intensively studied in the last thirty years and are still far from completely understood. Chaotic systems have been used to attempt to model a wide variety of real situations which exhibit unpredictable behavior: business cycles, turbulence, heart attacks, etc. Although fascinating and philosophically provocative, most of this work is still very speculative and has yet to prove of practical value.

Systems involving more than three variables can still be treated geometrically: we think of the space of states as a higher dimensional space (one dimension for each quantity) and the dynamical system as defining a vector field on the state space. Of course, we cannot visualize such spaces directly, but the geometrical insight we gain in dimensions two and three very frequently allows us to handle such systems.

Exercises

In the next two exercises, we look at some three-dimensional systems which arise in ecology. These questions are challenging and you will probably find it helpful to work them out in a group.

1. a) Consider a system consisting of three species: giant carnivorous reptiles, vegetarian mammals, and plants. Suppose that the populations of these are given by x, y and z respectively. The reptiles eat the mammals, the mammals eat the plants, and the plants compete among themselves. Explain why

the following system is consistent with these hypotheses:

$$\begin{aligned} x' &= -.2x + .0001xy, \\ y' &= -.05y - .001xy + .000001yz, \\ z' &= z - .00001z^2 - .0001yz. \end{aligned}$$

b) Find all equilibrium points of the system. There are five, one of which is physically impossible. Describe the significance of the other four.

c) The most interesting equilibrium is the one in which all three species are present. Localize the system at this equilibrium, using local variables u, v, and w. Linearize. Show that the linearized system has the form

$$\begin{aligned} u' &= \ .003v, \\ v' &= -2u + .002w, \\ w' &= -8v - .8w. \end{aligned}$$

Can you determine whether the equilibrium is an attractor? This is a hard question—it *is* an attractor. One way to show this is a generalization of the technique we used in the preceding section to examine the distance of points on a trajectory from the origin over time. For the current problem we use a **generalized distance function**

$$D(t) = 8 \cdot 10^6 u^2 + 12000v^2 + 3w^2.$$

Show, using arguments like those we used when we looked at ordinary distance, that as we move along a trajectory, the value of D must decrease. Hence conclude that the equilibrium point must be an attractor.

2. The system of equations

$$\begin{aligned} x' &= x - .001x^2 + .002xy - .1xz, \\ y' &= y - .01y^2 + .001xy, \\ z' &= -z + .001xz. \end{aligned}$$

arises in a general family of models proposed in 1980 by Heithaus, Culver, and Beattie (“Models of Some Ant-Plant Mutualisms,” *American Naturalist*, **116** (1980) pp. 347-361) for investigating the interactions three species: violets, ants, and mice. Violets produce seeds with density x (per square meter, say). The ants take some of the seeds and use the seed covering for food. But they leave the remainder, which is still a perfectly good seed, in their refuse piles, which happily turn out to be good sites for germination. The ants have density y. Finally, the seeds are also taken by mice, who use the whole seed for food (destroying both the cover and the seed within). The mice have density z.

a) Explain why these equations are consistent with the hypotheses we made on the interactions between the violets, ants and mice.

b) Find all equilibrium points for the system. Don't forget the points where one or more of the variables equals 0.

$$\begin{aligned} x' &= x - .001x^2 + .002xy - .1xz, \\ y' &= y - .01y^2 + .001xy, \\ z' &= -z + .001xz. \end{aligned}$$

c) Localize the model at each of these equilibria, using local coordinates u, v, and w as before, and linearize.

d) In the case of the equilibrium point (1000, 200, 4) the local linearization is

$$\begin{aligned} u' &= -u + 2v - 100w, \\ v' &= .2u - 2v, \\ w' &= .004u. \end{aligned}$$

As in the preceding problem, show that this point is an attractor by examining the generalized distance function $R(t) = u(t)^2 + v(t)^2 + 25000w(t)^2$ and showing that the value of R decreases as you move along a trajectory.

8.6 Chapter Summary

The Main Ideas

- A dynamical system can be viewed as a geometrical object. The possible values of the dependent variables are then the coordinates of a point—called a **state**. The set of all possible points is called the **state space** for the system.

- The differential equations become a rule assigning a **velocity vector** to each state. Thought of in this way, the equations are called a **vector field**.

- Solutions to the differential equations correspond to **trajectories** in the state space. At every point a trajectory is tangent to the corresponding velocity vector, and is changing at the rate given by the length of the vector. The set of all possible trajectories is called the **phase portrait** of the system.

- **Equilibrium points** are points where the velocity vector is 0. An equilibrium point is a trajectory consisting of a single point. A dynamical system is conveniently analyzed by examining the nature of its equilibrium points—whether they are **attractors**, **repellors**, **saddle points**, or **centers**.

- To study the nature of an equilibrium point it is helpful to look at the **local linearization** of the vector field near the point.

- Determining whether **fixed-line trajectories** exist is a crucial part of analyzing the nature of an equilibrium point.

- In addition to equilibrium points, dynamical systems in two dimensions may also have **limit cycles** that shape the asymptotic behavior of the system.

- In higher dimensional state spaces, there are not only the obvious extensions of point attractors and limit cycles, but it is possible to have an **attracting torus** as well. There are even more complicated attracting objects called **strange attractors**.

Expectations

- You should be able to describe the assumptions embodied in a particular dynamical system modeling the interaction between two (or three) species and evaluate whether the assumptions seem reasonable.

- For a dynamical system with two dependent variables, you should be able to determine the regions where each variable is zero or has a constant sign, find equilibrium points, sketch representative vectors of the vector field, and draw trajectories that are consistent with this information.

- You should be able to determine whether a linear system of differential equations with two dependent variables has **fixed-line** trajectories—that is, trajectories that are straight lines going directly toward or directly away from an equilibrium point.

- You should be able to **localize** and **linearize** a dynamical system in two variables to explore its behavior near an equilibrium point.

- You should be able to recognize the five generic types of equilibrium points: attracting and repelling **nodes**, attracting and repelling **spirals**, and **saddle points**.

- Using a differential equation solver, you should be able to recognize when a dnamical system has a **limit cycle**.

- You should be able to analyze a dynamical system with three dependent variables.