Chapter 1

Modeling, Change, and Simulation

1.1 Feedback

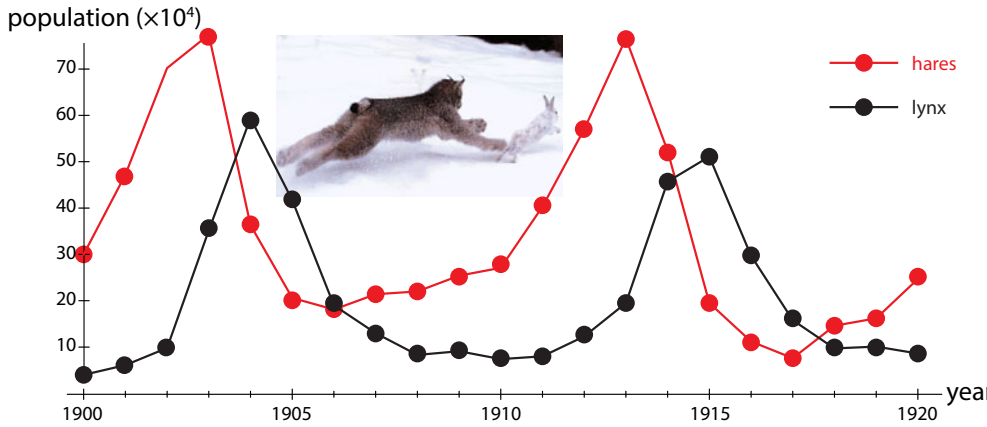

Figure 1.1: Oscillations in the populations of lynx and snowshoe hares over time, from pelts of animals captured by Hudson Bay Company trappers from 1900 to 1920.

In the 1920s, ecologists began to study the populations of two Arctic species, lynx (a predator) and snowshoe hares (their prey) (Figure 1.1). Notice that the populations *oscillate*. These oscillations are not random fluctuations; they have a roughly constant period of about 10 years. Note also that the rise and fall of the predator population systematically lags a little behind that of the prey population.

We want to understand what could be causing these oscillations. The first thing we have to realize is that finding the explanation requires us to take a careful look at the dynamic relationships between the two species. We have to make a *model* of their interactions. Even a very rough verbal model of the interaction reveals an interesting fact: the prey population *positively* affects the number of predators, while the predator population *negatively* affects the number of prey.

This makes the lynx–hare system our first example of a system with *negative feedback*.

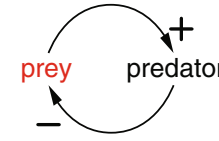

If we try to predict the system's behavior based on this verbal model, we discover a problem. Suppose we start with a certain number of predators

A. Garfinkel et al., *Modeling Life*,
DOI 10.1007/978-3-319-59731-7_1

and prey. What will happen? Well, the predators will eat some of the prey, and so the predator numbers will go up, and the prey numbers will go down, but then what? With high numbers of predators and low numbers of prey, the system cannot continue at the same pace; indeed, predator numbers will decline. But then what?

The problem here is that there is feedback in this system: the prey population affects the predator population and the predator population affects the prey population. It is difficult to predict the behavior of a feedback system based on this kind of verbal model. *The diagram above has to be turned into a real mathematical model if we want to predict and understand the behavior.*

The purpose of this book is to learn the art of making mathematical models of natural phenomena and learning how to predict behavior from them.

Interestingly, when we make a simple mathematical model of the predator–prey dynamics, it makes some nonobvious predictions.

This first model will leave out weather fluctuations, disease outbreaks, plant abundance, the effects of crowding on behavior, and innumerable other things. It will include only birth, death, and predation. To underline the fact that this is a highly simplified model, we will call the predators *sharks* and the prey *tuna*. You will learn more about this model in Section 1.4, but for now, we can look at the behavior that this highly simplified model predicts, shown in Figure 1.2. This kind of graph, which shows how quantities change over time, is called a *time series*.

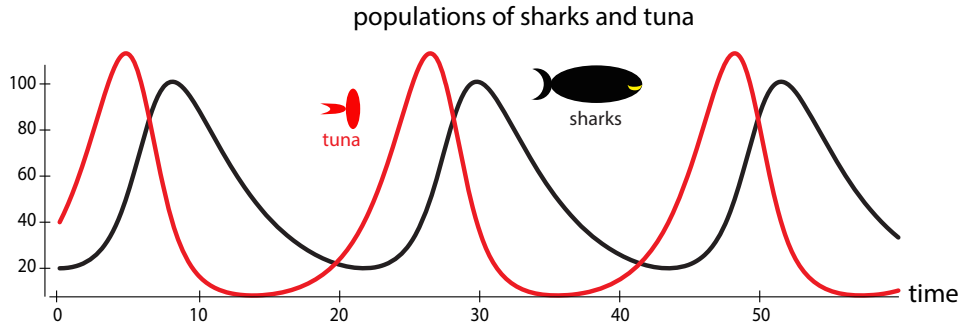

Figure 1.2: Behavior predicted by a model of interacting predator (shark) and prey (tuna) populations.

As you can see, the model predicts oscillations similar to the ones observed. What is the cause of these oscillations? The key to understanding this system is *time delays*. Sharks eat tuna, so the shark population grows and the tuna population diminishes until we get to a state where there are many sharks and very few tuna. The shark growth was caused by the previously high tuna levels, but now the tuna levels are low. The delayed shark growth has created a high-shark/low-tuna state, which means that the shark population will then decline, because due to the low tuna levels, very few sharks will be born and/or survive to maturity. This shark crash then takes the pressure off the tuna population, which then starts growing. The cycle then repeats.

Exercise 1.1.1 Copy two full cycles of the predator–prey oscillation time series in Figure 1.2 and label the point at which each of the processes described in the above paragraph is occurring.

This introductory example features oscillations of two species, hares and lynx, but it has been noted that these data are actually not populations of animals, but rather populations of *pelts*, as collected by hunters. Therefore, some ecologists have argued that a better model

for the populations would include *two* predators on the hares, lynx and hunters, each with their dynamics. This may well be true (Weinstein 1977). A model is a hypothesis about how to explain the data, and the merits of alternative models often must be considered.

Feedback Loops

The shark–tuna system is an example of a system with **feedback**. The tuna population has a positive effect on the shark population, while the shark population has a negative effect on the tuna population (Figure 1.3).

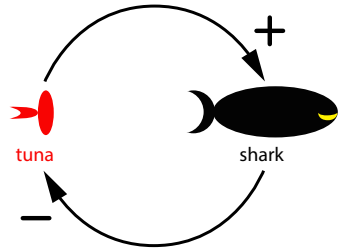

Figure 1.3: Shark and tuna feedback dynamics.

More generally, in a *feedback loop*, the current state of a system affects the future state of that system by changing the inflows or outflows of the system's components. There are two types of feedback loops: positive and negative.

Positive Feedback

In *positive feedback*, a positive value of a variable leads to an increase in that variable, and a negative value of a variable leads to a decrease (more negative) in that variable.

For example, a person who has money will be able to invest it, bringing in more money. Or think of practicing a sport or musical instrument. Practice makes you better at the activity, which makes you enjoy it more, which makes you practice more. Positive feedback can also be bad, which might be casually referred to as "negative." Think of a gambler who is losing badly, and so gambles more, ending up even further behind. In an arms race, one country can purchase more weapons, which causes its adversary to purchase more weapons, which causes the first country to purchase even more weapons. This is still positive feedback, just in a bad direction. Positive feedback reinforces change, so it can be thought of as "reinforcing feedback."

There are many important examples of positive feedback loops:

Population growth Animals have young, which increases the population. The larger the population is, the more babies are born, which makes the population even larger. As long as resources are available, the population will keep growing.

CO_2 emissions Carbon dioxide emissions trap heat, which raises global temperatures. At higher temperatures, soil microbes have faster metabolic rates, which means that they break down soil organic matter faster, releasing even more CO_2.

Methane release Methane is a greenhouse gas 25 times more potent than CO_2. Large amounts of methane are trapped in Arctic permafrost and at the bottom of the ocean. Rising temperatures cause this methane to be released, contributing to further temperature increases.

Market bubbles and crashes In a market bubble, investors buy into a stock, which causes the price to rise, which encourages more investors to buy, on the grounds that the stock is "going up." In a crash, investors sell the stock, which lowers the price, which convinces others to sell because the stock is "going down."

Negative Feedback

The other type of feedback is called *negative feedback*. In negative feedback, a positive value of a variable leads to a decrease in that variable, and a negative value of a variable leads to an increase in that variable. For example, a person whose bank account is low might work overtime to bring it back up and then cut back on overtime when there is sufficient money in the account.

This really is an example of negative feedback. We can often define a new variable from a given one by choosing a reference value for the variable and then defining the new variable as the given variable *minus* the reference value.

For instance, we can define B_0 as your desired bank balance; let's say $B_0 = \$1000$. Then if your current balance in dollars is D, we define a new variable B describing the discrepancy between D and B_0,

$$B = D - B_0$$

so a value of D that is less than B_0 will produce a negative value of B. This negative value will then increase (become less negative) under negative feedback.

The classic example of negative feedback is a thermostat that controls an air conditioner (Figure 1.4). When the temperature goes up, the air conditioner comes on, which causes the temperature to go down. To phrase this more carefully, we let T_0 be the set point on the thermostat. Then if the current air temperature is C degrees, we define the temperature T as

$$T = C - T_0$$

so that values of C above T_0 produce positive values of T, and values of C below T_0 produce negative values of T. The thermostat can also control a heater, in which case a decrease in temperature causes the heater to turn on, which raises the air temperature. In both cases, the thermostat opposes the change.

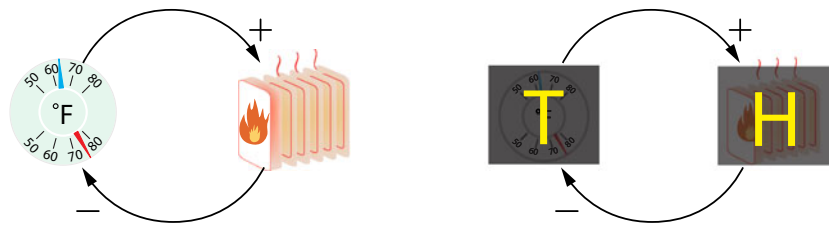

Figure 1.4: Feedback loop of thermostat and heater.

Insulin/Glucose Another example of a negative feedback loop involves insulin and glucose in the bloodstream. Intake of glucose (say, as a result of a meal) causes the pancreas to secrete more insulin, which then lowers the level of glucose by helping the glucose to be metabolized in the body. This feedback loop, which has time delays like those in the shark–tuna system, causes oscillations like those shown in Figure 1.5, even when a person is hooked up to a constant intravenous glucose supply with no meals.

Hormone regulation Virtually all of the hormones of the body are under negative feedback control by the brain and pituitary gland. For example, the gonadal hormones estradiol and progesterone (in females) and testosterone (in males) are under negative feedback regulation by the brain/pituitary system. This results in oscillatory behavior in many hormonal systems (Figure 1.6).

Gene regulation Many genes inhibit their own transcription, resulting in oscillating gene expression. For example, one protein that is essential in the early development of the embryo is called

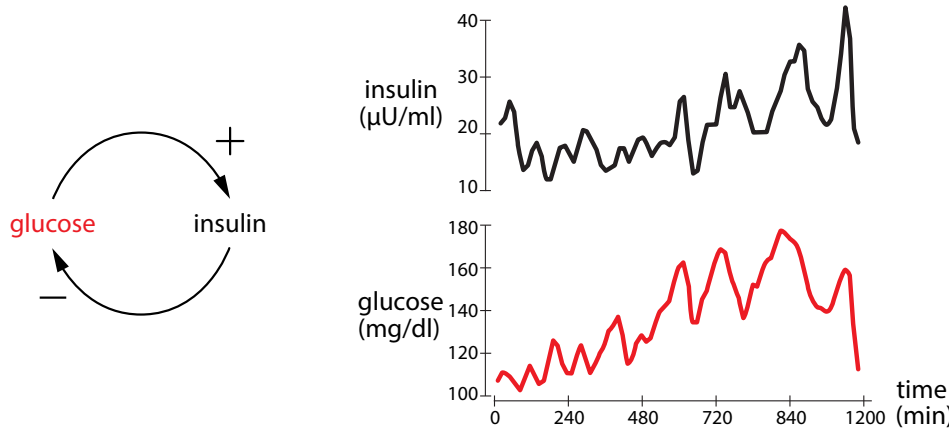

Figure 1.5: Feedback dynamics. Glucose and insulin concentrations in the blood of a person receiving a constant IV glucose infusion (Sturis et al. 1991a, b). Redrawn from "Aspects of oscillatory insulin secretion," by J. Sturis, K.S. Polonsky, J.D. Blackman, C. Knudsen, E. Mosekilde, and E. Van Cauter, In *Complexity, Chaos, and Biological Evolution*, by E. Mosekilde and L. Mosekilde, eds. (1991), volume 270, pp. 75–93. New York: Plenum Press. Copyright 1991 by Plenum Press. With permission of Springer.

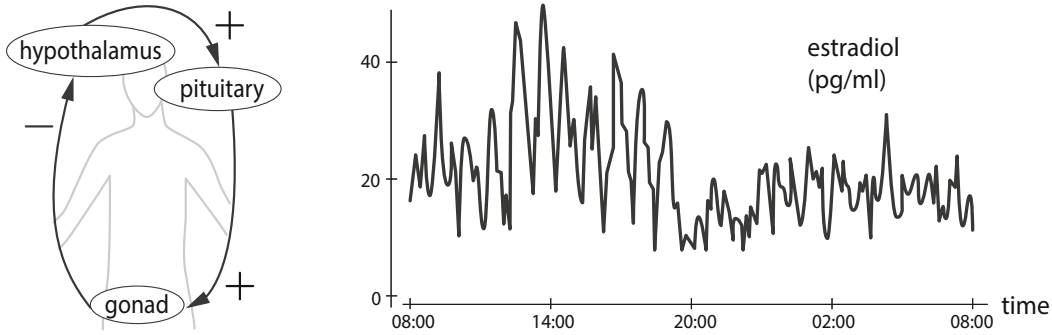

Figure 1.6: Left: the hypothalamus, pituitary, and gonads form a negative feedback loop, regulating the secretion of the sex hormones estradiol, progesterone, and testosterone. Right: oscillatory behavior in estradiol in a 24-year-old normal female, from Licinio et al. (1998). Redrawn with permission from "Synchronicity of frequently sampled, 24-h concentrations of circulating leptin, luteinizing hormone, and estradiol in healthy women," by J. Licinio, A.B. Negrão, C. Mantzoros, V. Kaklamani, M.-L. Wong, P.B. Bongiorno, A. Mulla, L. Cearnal, J.D. Veldhuis, and J.S. Flier (1998), *Proceedings of the National Academy of Sciences* 95(5):2541–2546. Copyright 1998 by National Academy of Sciences, U.S.A.

Hes1. Hes1 protein is produced by transcription from messenger RNA (mRNA). But then the protein inhibits its own transcription, producing a negative feedback loop. This leads to oscillations in protein levels (Figure 1.7).

Epidemiology Epidemiology is the study of diseases in populations. Many epidemiologists study infectious diseases. Contact between susceptible and infected people increases the transmission of the disease and causes the number of susceptible people to decrease. This decrease means that there are fewer susceptible people to infect, so transmission declines (see, for example, the epidemiology model on page 40).

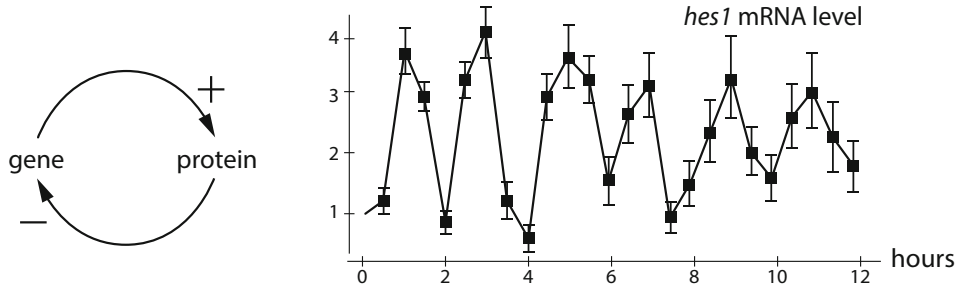

Figure 1.7: Left: many proteins inhibit their own genetic transcription, creating a negative feedback loop. Right: oscillations in mRNA levels of *Hes1* (Hirata et al. 2002). Redrawn from "Oscillatory expression of the bHLH factor Hes1 regulated by a negative feedback loop," by H. Hirata, S. Yoshiura, T. Ohtsuka, Y. Bessho, T. Harada, K. Yoshikawa, and R. Kageyama (2002), *Science* 298(5594):840–843. Reprinted with permission from AAAS.

During the Ebola epidemic in 2014, the U.S. Centers for Disease Control and Prevention (CDC) used a mathematical model of susceptible and infected populations to predict the course of the epidemic: how bad would it be? They also used the model to plan possible strategies for intervention: how much would we have to reduce the transmission rates to control the epidemic and even make the number of infected decline to zero?

Their results are shown in Figure 1.8. The left panel shows the predicted course of infection without intervention, while the right panel shows the effect of an intervention strategy that reduced the risk of transmission by getting 25% of patients into Ebola treatment units and 20% of the susceptible population into low-risk settings at first, and then gradually increasing that to 80% over six months. Note that this strategy is predicted to eliminate the epidemic.

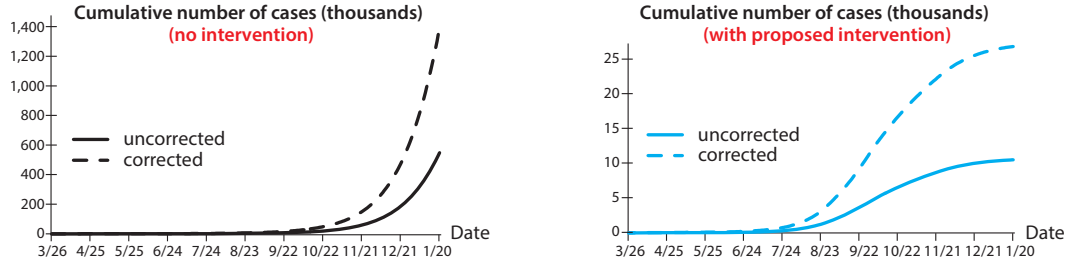

Figure 1.8: Predictions from the CDC Ebola transmission model. Solid lines show predicted reported cases, while the dashed lines show the predicted number of actual cases after correcting for underreporting (Meltzer et al. 2014). Note the numbers on the vertical axes.

Exercise 1.1.2 Come up with another example of a positive feedback loop and another example of a negative feedback loop.

Counterintuitive Behaviors of Feedback Systems

Most real systems consist of multiple feedback loops that interact. For example, a predator–prey system contains both a negative feedback loop, in which prey cause the predator population to increase and predators cause the prey population to decrease, and a positive feedback loop, in which a species causes its own population to increase through births.

For this and other reasons, systems with feedback often behave in counterintuitive ways. For example, suppose we want to reduce the number of sharks in an ecosystem. (This might actually be done in fisheries management.) We therefore remove sharks from the system. What happens?

The system responds by rebounding (Figure 1.9). Lowering the number of sharks takes the pressure off the tuna population, which grows to a higher level than before. The higher tuna population then gives rise to an even higher shark population. Thus, removing sharks dramatically actually results in a higher peak shark population!

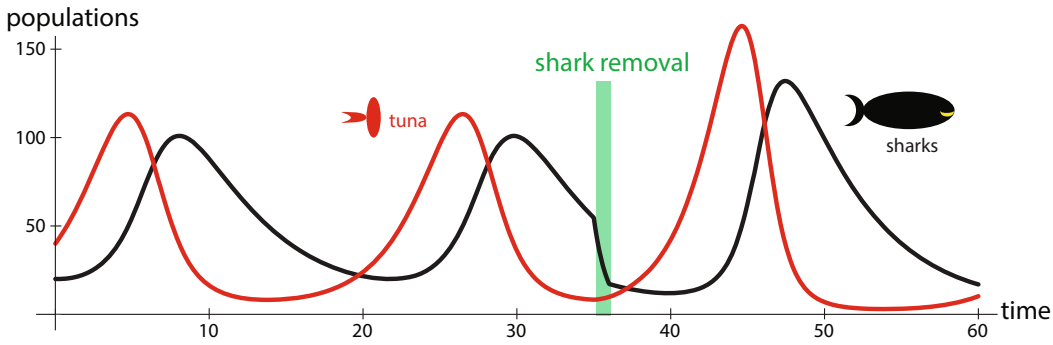

Figure 1.9: Complex systems defeat naive interventions.

In fact, the response of the feedback system to an intervention can depend strongly on the phase of the cycle in which the intervention is delivered, and can also depend on the magnitude of the intervention. There are interventions that decrease the shark population, and others that increase it.

In the simple model being simulated here, these high-amplitude oscillations after intervention continue indefinitely. This feature is a drawback of this model: it "remembers" perturbation forever. In Chapter 4, we will consider a better model that exhibits more robust oscillations (see *Stable oscillations in an ecological model* on page 200). The more advanced model also predicts the same counterintuitive response: the initial response of the system to a predator-removal intervention can be a rebound effect whereby the number of predators is increased, but transiently instead of permanently.

This basic principle, that feedback systems respond to intervention in counterintuitive ways, is seen throughout the natural world, from ecosystems to the human body. Testosterone is a hormone that enhances muscle building and is the drug that athletes most often abuse for performance enhancement. But testosterone, like all hormones, is under negative feedback control: sensors in the brain and pituitary gland register the amount of circulating testosterone and respond with negative feedback: they lower their output of testosterone-stimulating factors in response to higher levels of testosterone (see *The hypothalamic/pituitary/gonadal hormonal axis* on page 181). Consequently, when men use performance-enhancing drugs like testosterone and its analogues, the main symptom that is seen is testicular atrophy, caused by the shutdown of the native system due to the negative feedback.

Even simple systems with feedback can defeat naive intuition and frustrate naive intervention. We need to make models to keep track of behavior in such systems.

These examples should convince you that we need to learn how to model biological systems and predict their behavior. This is what we will now do. First, we need one crucial technical concept, the idea of a *function*, which you will learn about in the next section.

Further Exercise 1.1

1. At the start of a math class, some students do a little better than others because of better prior preparation, more time spent studying, etc. Students who do well feel confident and come to enjoy the class, leading them to spend more time on it. On the other hand, a student who does relatively poorly may decide that they're just not a math person and therefore put less effort into the class, thinking that it's not going to pay off. This, of course, leads to even lower grades, confirming the student's opinion.
   a) What kind of feedback loop is this?
   b) You are friends with a student who is having difficulty and losing confidence. How could you take advantage of the feedback loop to help your friend?

1.2 Functions

Think back to the graphs you saw when studying the shark–tuna predation model. At each point in time, the shark population has some value, and so does the tuna population. Look at Figure 1.10, which shows the tuna population.

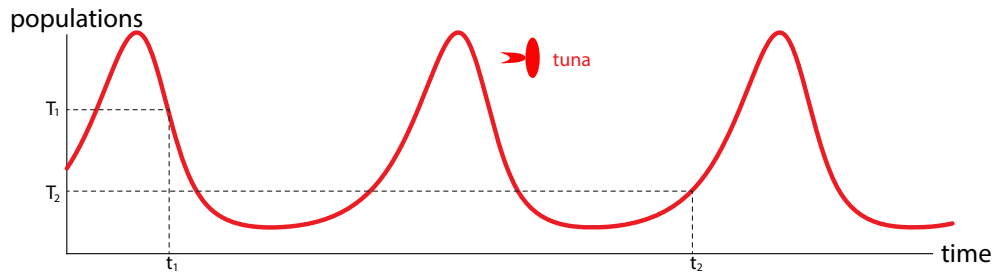

Figure 1.10: Tuna population in the shark–tuna dynamical system. Such a function of time is called a "time series."

Since there is exactly one population value at each time value, the tuna population is a *function of time*.

A *function* is a relationship between a set of inputs and a set of outputs, in which each input is assigned exactly one output—never none, never more than one.

One everyday example of a function is a menu that gives the prices of dishes at a coffee shop. Every drink has exactly one price—ordering would be rather confusing otherwise (Figure 1.11). A report card that gives a student's grades in different subjects is another example of a function.

HOT BEVERAGES	Price($)
Mocha	3.45
Cappuccino	3.45
Macchiato	2.45
Latte	3.15
Americano	3.75
Espresso	2.95

Figure 1.11: Like a coffee shop menu, a function associates one value (a drink) with another (its price).

Exercise 1.2.1 Come up with two more everyday examples of functions. Briefly explain what makes each example a function.

Functions can be thought of as machines, like the ones shown in Figure 1.12 and Figure 1.13. You put an input value into the function, and it returns a unique output value. The machine's behavior is absolutely predictable: the same input always produces the same output. This determinism is the defining feature of functions.

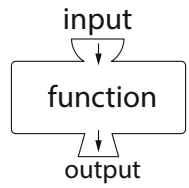

Figure 1.12: A function depicted as a machine.

Figure 1.13: A function, f, that takes either a blue star or a red star as input and switches the color.

Exercise 1.2.2 Modify the menu in Figure 1.11 so it no longer depicts a function.

Functions can also be defined by tables, with a little help. For example, one dataset recorded the amount of margarine consumption per person in the U.S. at various time points and also recorded the number of lawyers in New Jersey at those same time points. They provided a table of 10 such pairs of values (Table 1.1).

That table is defined for only 10 values. We can turn it into a function defined on the whole interval from the lowest margarine consumption to the highest using a technique like *linear interpolation*, which consists in simply drawing straight lines between your data points. The resulting function is shown in Figure 1.14.

Margarine consumption per person in the U.S. (lbs)	Lawyers in New Jersey
3.7	40,060
4.0	38,104
4.2	39,384
4.5	39,019
4.6	38,466
5.2	37,172
5.3	36,860
6.5	36,785
7.0	55,687
8.2	54,581

Table 1.1: 10 pairs consisting of margarine consumption per person in the US at various time points, together with the number of lawyers in New Jersey at the same time points. We have rearranged the 10 pairs in the order of increasing values of margarine consumption. Source http://tylervigen.com/view_correlation?id=29177

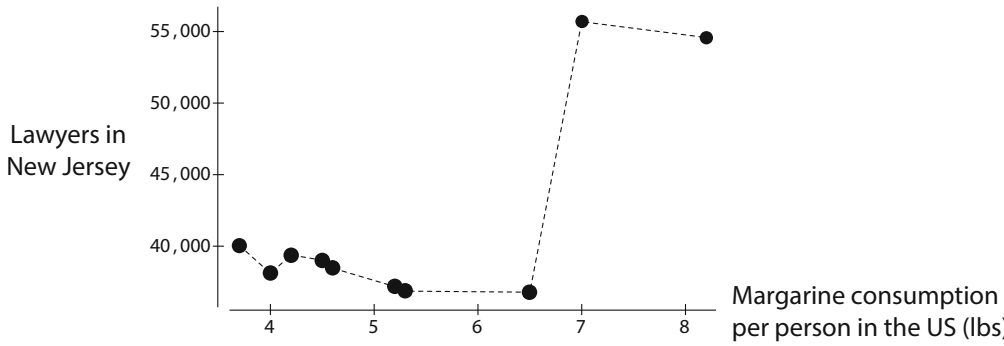

Figure 1.14: Linear interpolation between data points turns the table into a function defined for all numbers between 3.7 and 8.2, in which "Lawyers in New Jersey" is a function of "Margarine consumption per person in the US".

It's important to note that functions are not causal; there is no reason to think there is a causal relationship between the input and output of a function. The lawyers–margarine graph is an example of a function without a causal relationship.

In addition, despite what the machine metaphor implies, a function does not change one value into another any more than a menu changes foods into dollar values. A function is an *assignment* of one value to another.

Notation for Functions

Functions can be defined by tables, as in the coffee shop menu and lawyers examples, but most of the time they will be defined by formulas. For example, the function X^2 can be thought of as a machine that takes an input X and returns the value X^2.

The output that corresponds to a particular input to a function is written as

$$\text{FunctionName(input)} = \text{output}$$

A common name for functions is f, so we might write $f(X) = Y$. For example, the left half of Figure 1.15 gives two input–output pairs that define a function f. We can write this function as $f(3) = 5$ and $f(4) = 6$. (The symbolic expression $f(3)$ is pronounced "f of 3.")

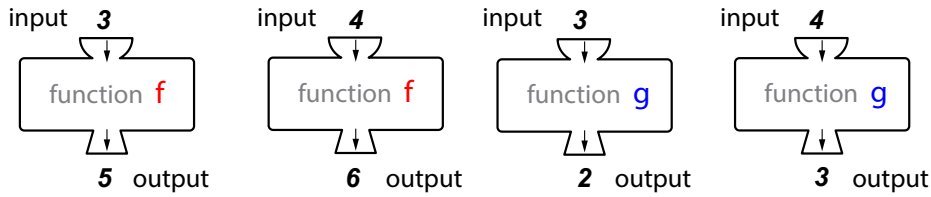

Figure 1.15: Two examples of very simple functions. The left half defines the function f, and right half defines the function g. Both functions are defined on the domain consisting of the two numbers 3 and 4.

Exercise 1.2.3 Use the two input–output pairs on the right side of Figure 1.15 to write the function g.

Writing functions using input–output pairs or tables of values works only if the number of values we're working with is relatively small. The compilation and use of such lists quickly becomes impractical as the number of values increases. Worse, many of the functions we frequently encounter have an infinity of possible values. Clearly, we need a better way of representing functions.

Sometimes, we can summarize the relationship represented by a function as a formula. For example, if $f(1) = 2$, $f(2) = 4$, $f(3) = 6$ and so on, we can write down the function by giving its name, the input variable, and then the expression that lets us find the value of the output given the input. In this case, the function is written as $f(X) = 2X$. Another simple example is $h(X) = X + 1$. The function is named h, the input variable is X, and the expression relating the output variable to the input variable is $X + 1$.[1]

Exercise 1.2.4 Write the functions f and g in Figure 1.15 in function notation using formulas.

[1] It's common to refer to input variables as *independent variables* and output variables as *dependent variables*.

It's important to realize that not every equation generates a function. Think about $X^2 = Y^2$. For the one value $X = 2$, we have two Y values: $Y = 2$ and $Y = -2$. So this equation does not define a function.

Interestingly, many (in fact, most) functions cannot be written as equations. Look back at Figure 1.10 on page 8. It's clearly the graph of a function, and a rather simple-looking one at that—maybe something like a sine or cosine function. Actually, though, *there is no known formula for the graph in* Figure 1.10.

Let that sink in for a bit. We have a simple-looking graph produced by a model that's not very complicated using a method you will learn about soon. And no equation or formula is known that gives us this graph!

This is not an exception. Actually, it's more like the rule. For the overwhelming majority of biological models, there is no known formula for the time series. However, an understanding of functions will prove very useful in studying these models.

Inputs and Outputs

The set of input values that a function can accept is called its *domain* of the function. In many cases, when a function is given by a formula, its domain consists of all real numbers. *Real numbers* are every kind of number you can think of as representing a quantity, including whole numbers (0, 1, 2, ...), fractions ($\frac{7}{3}$, etc.), irrational numbers ($\sqrt{2}$, π, etc.), as well as the negative values of all of these ($-\sqrt{2}$, etc.). Altogether, these numbers make up the real numbers, abbreviated $\mathbb{R}$ and pronounced "r." For convenience, we will define $\mathbb{R}_+$ (pronounced "r plus") to mean the nonnegative real numbers: 0 and everything larger.

The domain of a function is something we decide. While issues like division by zero restrict our choices in some ways, we are generally free to define the domain however we want. For example, the domain of $f(x) = \frac{1}{X}$ cannot include 0, because division by zero is undefined, but otherwise, the domain of $f(x) = \frac{1}{X}$ could consist of all real numbers greater than zero, or all integers (positive and negative) except zero, or even just the interval $[3, 7]$. In modeling real-world systems, it will be important to pick domains that make physical sense.

Exercise 1.2.5 Give three possible domains for a function defined by the formula $g(X) = \frac{2}{X-5}$.

There is also a term for a set of values in which a function's output lies. This is called the *codomain*[2] of the function. For example, the codomain of $g(X) = X^3$ consists of all real numbers. A function links each element in its domain to some element in its codomain. Each domain element is linked to exactly one codomain element. This is what makes functions unambiguous.

In many situations, it is useful to specify the domain and codomain of a function even if we don't specify the actual rule or formula by which domain elements are associated with codomain elements. In these cases, we often describe the function using the notation

$$\text{function name} : \text{domain} \to \text{codomain}$$

[2] In high school, you probably encountered the term "range" rather than "codomain." However, "codomain" is the accepted term in more advanced work. Technically, the range consists of the outputs a function actually gives, so finding the range of something like $f(X) = 2\sin X - 5$ takes a little calculation. The codomain, on the other hand, is just a set that includes all the values the function could return, so the codomain of f in this example can be said to be all real numbers.

For example, the coffee shop menu in Figure 1.11 on page 9 links drinks to prices. Therefore, we might describe the menu as a function by writing

$$\text{menu} : \{\text{drinks}\} \to \{\text{prices}\}$$

(Curly braces are standard mathematical notation for sets, so here, for example, we are using {drinks} to denote the set of drinks served by the coffee shop—the domain of this function.)

The notation $f : X \to Y$ is pronounced "f takes X to Y."

It's important to distinguish between the entire domain of a function and an element of the domain. In the menu example, the domain consists of *all* the drinks on the menu. A single drink is an element of the domain.

Exercise 1.2.6 Describe the everyday function examples you came up with in Exercise 1.2.1 on page 9 in "function name : domain → codomain" notation.

Putting Functions Together

An interesting example of functions comes from molecular biology. DNA encodes information about the makeup of proteins in a sequence of four base pairs, A, C, G, and T. When DNA is transcribed into RNA, T is replaced by another base, abbreviated U. Thus, transcription is a function in which A, C, and G are associated with themselves, but T is associated with U. We can write

$$\text{transcription} : \{A, C, G, T\} \to \{A, C, G, U\}$$

Things get more interesting when we consider not single bases but base triplets called codons. In transcription, each DNA codon is transcribed into an RNA codon. Then, each RNA codon causes a particular amino acid to be added to a protein. (There are also codons that start and stop the protein-building process, but we can ignore those for now.) This process is called translation and is also a function, because a particular RNA codon unambiguously specifies an amino acid. Thus, we have two functions:

$$\text{transcription} : \text{DNA codons} \to \text{RNA codons} \qquad \text{and}$$
$$\text{translation} : \text{RNA codons} \to \text{amino acids}$$

What about the overall process of gene expression, in which DNA codes for a protein? We can write down another function:

$$\text{gene expression} : \text{DNA codons} \to \text{amino acids}$$

This function is made up of the previous two functions linked end to end. The output (or codomain) of transcription becomes the input (or domain) of translation.

This kind of linking is called *composition of functions* and is the most natural way to combine functions. For example, if $f(X) = 2X + 1$ and $g(Y) = \sqrt{Y}$, then $g(f(X)) = \sqrt{2X + 1}$ and $f(g(Y)) = 2\sqrt{Y} + 1$. Composition of functions will become important later in this course.

Exercise 1.2.7 Suppose life is discovered on Mars. The Martians' genetic code is remarkably similar to ours, but the RNA codon AUC is translated to serine 60% of the time and histidine 40% of the time. Is Martian gene expression a function?

Exercise 1.2.8 As we saw earlier, a coffee shop menu is a function. Suppose that when you buy a drink, you have to pay 10% sales tax in addition to the price of the drink, so the total cost (price and tax) of a drink is 1.1 times the price on the menu.

a) Refer to Figure 1.11 on page 9. What is the total cost of a mocha? A latte?

b) Describe the process of finding the total cost in terms of function composition.

Further Exercises 1.2

1. Consider the restaurant menu below:

Item	Price
Pizza slice	$2.50
Hamburger	$4.00
Cheeseburger	$4.50
Fries	$2.00
Kale foam on a bed of arugula	$37.50

   a) Is price a function of item? Justify your answer.

   b) Is item a function of price? Justify your answer.

   c) Create a similar menu in which price is not a function of item and explain why it is not a function.

2. In high school, you may have learned the *vertical line test* for checking whether a graph is the graph of a function. (A graph is the graph of a function if a vertical line drawn through the graph intersects it exactly once, no matter where the line is drawn.) Explain why the vertical line test works.

3. Some ten-year-olds are experimenting with secret codes. Aisha's favorite code involves replacing each letter with the one two letters later, so A is replaced with C, B with D, Y with A, and Z with B. Meanwhile, Tim prefers to replace letters with their position in the alphabet: 1 for A, 2 for B, and so on. Suppose Aisha encodes the word "spam" with her code and Tim encodes the result with his code.

   a) What will the outcome be?

   b) Describe this scenario in terms of functions and their composition.

   c) Suppose Aisha's code is defined only for letters, not numbers. Could the kids apply Tim's code and then Aisha's code to a message? Explain your answer in terms of domains and codomains.

d) What does this example tell you about function composition?

4. What's wrong with the "function" $f(X) = \log(\log(\sin X))$? Your answer should involve domains and codomains. *(Hint: Try plotting $f(X)$ by hand.)*

5. A DNA codon codes for exactly one amino acid, but there are amino acids that are coded by several different codons. Is there a function that takes amino acids to DNA codons? Justify your answer.

6. In high school, you may have learned about function composition as just another way of putting functions together, similar to addition and multiplication. How is composing functions different from adding or multiplying them? *(Hint: Think about when we can do one but not the other.)*

7. In SageMath, let $a = 5$. Apply some SageMath function to a and view the result. Then, view a. Did its value change? Do this for two more values of a, using a different function each time and viewing a after *each* computation. What does this tell you about functions?

8. SageMath has a useful command, `simplify`, that algebraically simplifies symbolic expressions. What happens if you enter `simplify( (x^2)^(1/2) )` into SageMath? Now type in `assume(x >= 0)` and try again. What happens? What accounts for the difference? *(Hint: Think about domains.)* When you finish this problem, execute the command `forget()` to stop forcing SageMath to assume that `x` is nonnegative.

1.3 States and State Spaces

The State of a System

One of the key ideas we will use throughout this course is that of the **state** of a system. *State* is just a term for the condition of the system at a given time. (Think "state of the union.") For example, the state of your bank account might be the amount of money in it, the state of a traffic signal may be its current light color, and the state of a population of animals or cells might be its size.

We can discuss the state of a system only after we've decided what variables to focus on. For example, a color-blind person might describe the state of a traffic signal by the position of the light rather than its color, while an electrical engineer might completely ignore which light is on and focus instead on the internal workings of the traffic signal. Similarly, in describing an animal population, we might be interested in its sex ratio, the numbers of animals in different age classes, the distribution of individuals in space, or all of the above. The variables we use to describe the state of a system are called *state variables*.

Our choice of state variables is determined by both the structure of the system (are animals distributed more or less evenly over the landscape, or are there distinct subpopulations?) and how we plan to use our model (Figure 1.16). **Deciding what variables to focus on is often one of the hardest parts of building a model.**

When making mathematical models, we have to describe the state of the system using a number or a list of numbers. In this course, state variables must be *quantitative*. A state variable is therefore a quantity that describes some property of the system, such as its velocity, shark population, or blood glucose concentration.

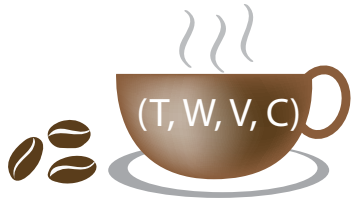

Figure 1.16: Four possible state variables defining the state of a cup of coffee.

State variables have units. For example, velocity might be measured in "meters per second" ($\frac{\text{m}}{\text{s}}$) or "miles per hour" ($\frac{\text{mi}}{\text{h}}$), a shark population is measured as number of individuals, and blood glucose concentration is commonly measured in "milligrams per deciliter" ($\frac{\text{mg}}{\text{dL}}$).

Exercise 1.3.1 Give possible units for measuring the following variables. Feel free to look up information as necessary.

a) Population density of prairie dogs

b) Concentration of epinephrine in the bloodstream

c) Amount of energy in a battery

The state of the system at any given time is given by the values of all of its state variables, in the appropriate units. For example, we might say that the state of a person right now is that their core body temperature is 101 °F, if body temperature is all we are interested in. A fuller description might be that the person's temperature is 101 °F, their systolic blood pressure is 110 mmHg, and their heart rate is 85 beats per minute.

As a system changes over time, the values of the state variables will change. Since a state variable can have only one value at a given time, *the values of state variables are really functions of time*. For each point t in time, we have a value of X. When we refer to X (e.g., "sharks"), we really mean X at a time t. So the state variable X is really a function of time. However, while X is really $X(t)$, we will usually just write it as "X" and leave the "(t)" implicit.

Since a state variable is a function of time, we can plot this function as a graph, with time on the horizontal axis and the state value on the vertical axis. This is an extremely important kind of plot called a *time series*. Tuna population in the shark–tuna dynamical system in Figure 1.10 on page 8 is an example of a time series.

State Space

When we work with dynamical models, our primary interest is not in learning what state a system happens to be in at a particular time. Rather, we want to understand the system's *behavior*—its changes from state to state—and why it exhibits one pattern of behavior rather than another. For example, we want to know why hare and lynx populations in Canada undergo multiyear cycles rather than remaining at roughly the same value each year or changing in a much more unpredictable way. In order to do this, we have to consider all possible system states and then ask why the system moves among them in particular ways.

The set of all conceivable state values of a system is called its state space—literally, the space in which the system's state value can move. For example, the state space for the color of a traffic light is {red, yellow, green}. Similarly, the state space for the behavior of a cat might

be {eating, sleeping, playing, walking on your keyboard}. But in this course, every state will be a *number*, such as temperature, number of animals, or glucose concentration.

The assumption of continuity

The state spaces we deal with in dynamics typically are spaces whose state variables are continuous. In other words, while 3457 is a valid value of X, so are 3457.1, 3457.12437, and even a number whose decimal expansion goes on forever without repeating, such as $\sqrt{2}$. We make the assumption of continuity even when our state variable is the size of a rabbit population. We don't worry too much about what it means to have 3457.1 rabbits. The same thing happens in chemistry—the number of molecules of a compound in a one-liter solution must be an integer, but it's such a large integer that we approximate it with a real number. Usually, this works well and allows us to use powerful mathematical tools. However, when you get down to the case that there are only three rabbits in your population (as sometimes happens in conservation biology and other fields), the assumption of continuity goes badly wrong, and you need to move to a different kind of modeling. Similarly, in chemistry, when your beaker has only three sodium ions, you also need to adopt a different kind of modeling.

One-Variable Systems

If a system has only one variable, its state is a real number (which in a given case might be restricted to being nonnegative). Thus, we can use the fundamental idea that the real numbers exactly correspond to points on a line to say that *the state space of such a system is a line*.

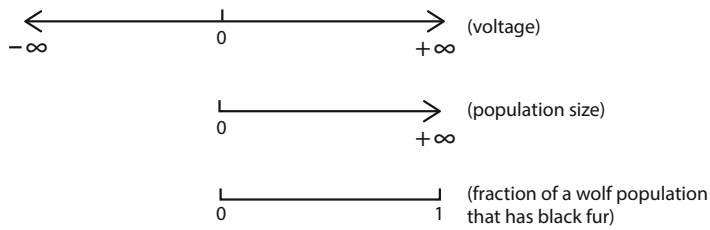

Figure 1.17: Three examples of one-dimensional state spaces

A system's state is represented as a point in state space, which we will sometimes call the *state point*. For a variable that can be either positive or negative, such as voltage, the state space is just the real number line. For a population size, the line goes from zero to infinity, excluding negative values. (Of course, we draw only the portion of interest.) For a proportion, like the fraction of a wolf population that has black fur, the line goes from zero to one (Figure 1.17).

Exercise 1.3.2 What is the state space for the number of ants in an ant colony?

Exercise 1.3.3 What is the state space for temperature measured in degrees Celsius? (Be careful!)

State Spaces with Multiple Variables

So far, we've seen a state space with one variable. That state space is a line, which can be thought of as a one-dimensional space. It may be a bit counterintuitive to think of a line as a "space," but that's just because we are used to thinking of the 3-dimensional space that we live in. A line really is a 1-dimensional space—a point can move on a line. But now we will go on to use more than one state variable, and the state spaces will start to look more like spaces. Here, the idea of state space really comes into its own.

Think of the shark–tuna system. We need two numbers to describe its state at a particular time, namely, the size of the shark population, S, and the size of the tuna population, T. Then the state of the shark–tuna system is given by a pair of numbers (S, T), which we write in parentheses with a comma between them. Note that order counts: $(3, 6)$ is not the same as $(6, 3)$. A system with 3 sharks and 6 tuna is different from a system with 6 sharks and 3 tuna.

Doing Math with States

We can work with states mathematically. Starting with the one-variable case, we can define two simple operations on states:

- If X_1 and X_2 are two values of the state variable X, then we can *add* them to produce another value of X:
$$X_3 = X_1 + X_2$$
We can do this because we know how to add two real numbers. We can always add apples to apples and sharks to sharks. For example, 3 volts + 5 volts = 8 volts.
- If X_1 is any state value, we can always multiply that state value by a number. Such a number is called a *scalar*. We can do this because we know how to multiply two real numbers. If X_1 is a state value, then so are $2.5X_1$, πX_1, etc. So, for example, 3(5 volts) = 15 volts.

Of course, we should perform such operations only when they make physical sense. Multiplying a population size by a negative scalar would give you a negative population. If we are talking about raw population numbers, then this is physical nonsense.

The rules for operating with pairs of values are similar. We just take into account the fact that we can add apples to apples and distances to distances, but not apples to distances.

- Pairs can be added. In the shark–tuna system, if (S_1, T_1) and (S_2, T_2) are states, then since we know how to add S's and how to add T's, we can define
$$(S_1, T_1) + (S_2, T_2) = (S_1 + S_2, T_1 + T_2)$$
If the state space is (apples, oranges), then we add apples to apples and oranges to oranges.
- Pairs can be multiplied by a scalar. If a is a scalar and (S_1, T_1) is a state, then since we know how to multiply S and T by scalars, we can define
$$a(S_1, T_1) = (aS_1, aT_1)$$

For example,

$$3.5(2 \text{ apples}, 3 \text{ oranges}) = (7 \text{ apples}, 10.5 \text{ oranges})$$

Exercise 1.3.4 Compute the following:

a) $5(10, 2)$ b) $(4, 7) + (3, 9)$ c) $2(3, 2) - 3(5, 4)$

The Geometry of States

If one number corresponds to a point on the one-dimensional line, what does a pair of numbers correspond to? For sharks and tuna, we can draw one line for the S variable and another line for the T variable. We can then make those lines perpendicular to each other and think of them as the axes for a two-dimensional space, called "shark–tuna space," as shown in Figure 1.18.

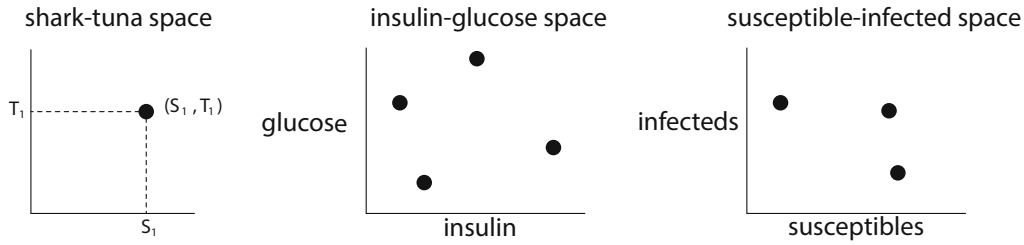

Figure 1.18: Examples of two-dimensional state spaces

A system's state space is often named by its variables. For example, the state space whose variables are insulin and glucose concentrations is called "insulin–glucose space" and that of a model of susceptible and infected populations can be referred to as "susceptible–infected space." A point in insulin–glucose space represents a particular concentration of insulin combined with a particular concentration of glucose; a point in susceptible–infected space represents a particular susceptible population size and a particular infected population size.

Exercise 1.3.5 Suppose we are modeling a black-bear population consisting of juveniles and adults. Draw the appropriate axes and a point representing the state of the black-bear population if there are

a) 200 juveniles and 100 adults

b) 30 juveniles and 50 adults

c) 0 juveniles and 25 adults

Exercise 1.3.6 Pick a two-variable system of any kind and draw its state space and a point representing a system state. Describe the state this point represents.

The concept of "shark–tuna space" is critical in this course. Don't confuse this with the physical space that the sharks and tuna swim around in; this is different. This is an abstract space whose coordinates are not physical positions but "number of sharks" and "number of tuna."[3]

Generalizing, if X and Y are state variables, then the set of pairs (X, Y) is the set of all states of the two-variable system. Since typically, X and Y will both be real numbers drawn from $\mathbb{R}$, we call the space of all pairs of real numbers $\mathbb{R} \times \mathbb{R}$ (pronounced "R cross R") or $\mathbb{R}^2$ (pronounced "R two" or "R squared").

We will now introduce some terminology. A fancy name for a pair of numbers is a 2-*vector*. The numbers making up the vector are called its *components*. The space $\mathbb{R}^2$ is called a two-dimensional *vector space*.[4]

This definition of "vector" may be slightly new to you. You may remember from high school that we can view vectors as arrows. Here, vectors are points in a vector space. The two views of "vector" can be reconciled by letting the vector $(3, 7)$ represent $(S = 3, T = 7)$ or as the arrow from $(0, 0)$ to $(3, 7)$, as in Figure 1.19. We will use both representations of vectors heavily.

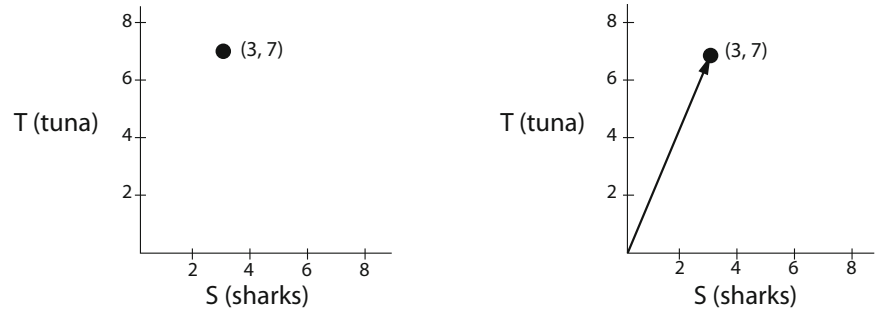

Figure 1.19: Left: vector as point. Right: vector as arrow from $(0, 0)$. Both representations carry the same information.

Scalar multiplication of vectors has an important geometric interpretation, shown in Figure 1.20 on the next page. Multiplying a vector by a scalar stretches the vector if the absolute value of the scalar is greater than 1, and it compresses it if the absolute value of the scalar is less than 1. If the scalar is positive, then the result is a vector in the same direction. What about a negative number? Numerically, multiplying a vector by a negative number changes the signs of all of the vector's components. Geometrically, this flips the direction of the vector, in addition to stretching it by the absolute value of the number.

Similarly, the addition of two vectors can be represented geometrically, as shown in Figure 1.21. If we add the vector $(8, 4)$ to the vector $(1, 3)$, the algebra of vectors tells us that $(8, 4)+(1, 3) = (9, 7)$.

The figure makes it clear that the vector $(9, 7)$, if we think of it as an arrow, is what we would get if we could put the base of the arrow representing $(8, 4)$ right on the tip of the arrow

[3] The idea of such an abstract space was developed by the mathematician Bernhard Riemann (1826–1866), who spoke of "multiply extended magnitudes" and said that "physical space is only a particular case of a triply extended magnitude" (Riemann 1873).

[4] Technically, only state spaces in which all variables can be both positive and negative are vector spaces, but this does not affect anything we do in this book.

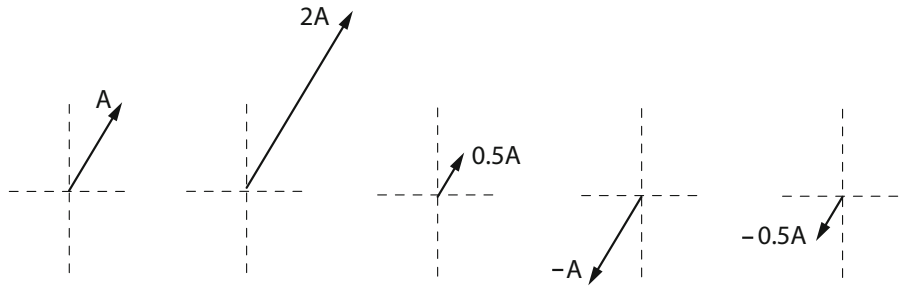

Figure 1.20: The result of multiplying the vector A $(1, 1)$ by 2, 0.5, -1, and -0.5.

representing $(1, 3)$, and thereby "adding" $(8, 4)$ to $(1, 3)$. Notice that the reverse procedure, adding the arrow $(1, 3)$ to the arrow $(8, 4)$, gives the same answer.

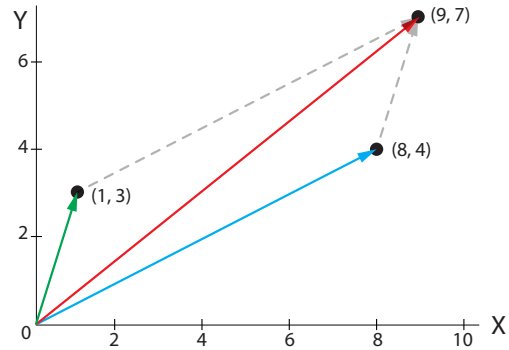

Figure 1.21: Vector addition. The red vector is the sum of the blue and green vectors.

Exercise 1.3.7 Draw two vectors and the same vectors multiplied by -1.

Exercise 1.3.8 Draw two vectors and show their sum.

State Spaces with More than Two Dimensions

You already know that a single number gives the position of a point on a line, and an ordered pair specifies the position of a point in a plane. Similarly, if a model tracks the concentrations of three different chemical compounds, its state space is a three-dimensional space whose axes represent the concentrations of the compounds in question. An ordered triple specifies the position of a point in three-dimensional space. Generalizing this idea, a vector with n components gives the position of a point in n-dimensional space.

Since the shark–tuna model has two variables, we need only two axes to specify its state. For more variables, we need more axes—one axis per variable. The number of axes needed to represent a system's state is called the *dimension* of its state space. In this text, we will pay particular attention to two- and three-dimensional models, because we can easily visualize their

state spaces, but most models used in research are much larger. Therefore, *we need to learn to work with vectors with any number of components.*

Operations on n-vectors are straightforward generalizations of those on 2-vectors. Vector addition is done componentwise: if $\mathbf{a} = (a_1, \ldots, a_n)$ and $\mathbf{b} = (b_1, \ldots, b_n)$, then

$$\mathbf{a} + \mathbf{b} = (a_1 + b_1, \ldots, a_n + b_n)$$

(When we want to talk about a whole vector without listing its components, we write its name in **boldface**.) Vectors can be added only if they have the same number of components.

Multiplying a vector by a scalar is also straightforward. Suppose a bear population has 100 juveniles and 200 adults. We triple our sampling area and find that the ratio of adults to juveniles is the same in the larger area as in the smaller one. To obtain the numbers of juveniles and adults in the larger area, we just multiply the numbers from the smaller area by 3. In vector notation, $3(100, 200) = (300, 600)$, and more generally,

$$c(a_1, \ldots, a_n) = (ca_1, \ldots, ca_n)$$

Unfortunately, we can't show you a picture of 4-dimensional or 50-dimensional space. We lowly humans cannot visualize four dimensions, let alone 11 or 27. But this is no problem! We can't draw or visualize 27-dimensional space, but if we need 27 variables to describe the state of a system, we just form the 27-vector $(x_1, x_2, x_3, \ldots, x_{27})$ and operate on it with the rules of vector addition and scalar multiplication as defined earlier.

Exercise 1.3.9 Carry out the following operations, or say why they're impossible.

a) $(1, 2, 3) + (-2, 0, 5)$

b) $-3(4, 6, -9)$

c) $(2, 4) + (1, 3, 5)$

d) $5\Big((0, 1) + (7, 3)\Big)$

Previously, we used the symbol $\mathbb{R}$ to refer to the real number line, and the symbol $\mathbb{R}^2$ to refer to two-dimensional space. Extending this idea, we can think of an n-dimensional space as having n copies of $\mathbb{R}$ as axes and denote it by $\mathbb{R}^n$ (pronounced "R n").

$$\mathbb{R}^n = \underbrace{\mathbb{R} \times \cdots \times \mathbb{R}}_{n \text{ times}} = \{(x_1, \ldots, x_n)\}$$

Exercise 1.3.10 How would you symbolize a 3-dimensional state space in this notation? an 18-dimensional state space?

Change

In this geometric picture, what is change? **Change is movement through state space** (Figure 1.22).

When a system changes, its state changes. In the figure, the system has changed from $x = 4$ at time t_1 to $x = 6$ at time t_2. The same idea that change is movement in state space also applies in higher-dimensional spaces. For example, if a predator–prey system goes from having,

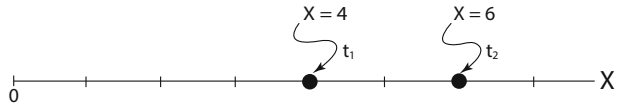

Figure 1.22: The state of the system at a time is given by a point in state space.

say, 3 tuna and 7 sharks to having 5 tuna and 10 sharks, its state changes from $(3, 7)$ to $(5, 10)$. Geometrically, this means that the state point moves in state space, from the point $(3, 7)$ to the point $(5, 10)$.

This is a powerful idea that will serve us throughout this course. We will now take up the question of what makes a state point move, i.e., the causes of change.

Further Exercises 1.3

1. This section defined vector addition and multiplication by scalars. Use these operations to compute $\begin{pmatrix} 5 \\ 1 \end{pmatrix} - \begin{pmatrix} 3 \\ 2 \end{pmatrix}$, justifying each step.
2. (From Bodine et al. (2014).) A state park consists of 80 acres of meadow, 400 acres of pine forest, and 520 acres of broadleaf forest. The park has the opportunity to acquire a parcel of land consisting of 25 acres of meadow, 130 acres of pine forest, and 300 acres of broadleaf forest. Write this as a sum of vectors and find out how many acres of each ecosystem type the enlarged park would consist of.

1.4 Modeling Change

Change is movement through state space. Now we want to go beyond this *description* of change, to talk about the *causes* of change. A set of hypotheses about the causes of change in a given system is called a *model*.

A Simple Example

Let's start with a simple situation: the amount of water in a bathtub.

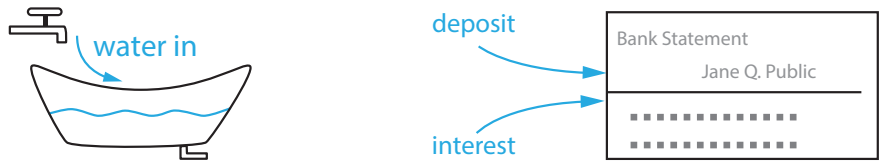

Figure 1.23: Two examples of systems with a single state variable and inflows that increase the value of that state variable. The level of water in the bathtub is increased by the flow from the faucet, and the bank balance is increased by the flows from deposits and interest.

Let's describe the state of the bathtub as

$$X = \text{amount of water in the tub (in gallons)}$$

What is changing the amount of water in the bathtub in Figure 1.23? The faucet, or to be more precise, the inflow of water through the faucet.

The units of this flow are *not* gallons, but gallons *per minute* (or some other time unit). We write that as $\frac{\text{gal}}{\text{min}}$. It's not "stuff"; it's "stuff per unit time."

The idea is that **levels** *are changed by* **flows**; *that is, quantities are changed by rates.* Your bank account balance (a quantity of, say, dollars) is changed by your income (in, say, dollars per month) and your expenses (also in dollars per month).

We represent this by a "change equation," in which we take the state variable X and define X' ("X prime") as the change in X. Then we write

$$X' = [\text{the things that change } X]$$

For example, for the bathtub above, we would write

$$X' = \text{faucet}$$

Now of course we haven't specified "faucet" yet, but we know that it has to be in gallons per minute. Let's make the assumption that the flow is constant over time, and that its value is 10 $\frac{gal}{min}$. We then write

$$X' = 10$$

This is our first example of a change equation, or model, with no words, just mathematical symbols representing the various causes of change.

Similarly, if X is your bank account balance and you never withdraw money, then a change equation for the account balance would be $X' = D + I$, where D represents your deposits and I represents interest, both in dollars per month (or year).

Let's go on to another example, with negative terms (Figure 1.24).

Figure 1.24: On the left, the drain in the bathtub provides an outflow that reduces the level of water. On the right, monthly withdrawals for rent, etc. reduce the level of the bank balance.

Now our change equation is clearly going to be

$$X' = -\ \text{outflow}$$

Note the minus sign. The outflow is clearly going to subtract from X, and make its value less, so it has to have a minus sign. But what is the "outflow"?

Now we have a situation we haven't seen before: the flow out of the bathtub is not constant; it **depends** on the amount of water in the bathtub. This is our first case of feedback. The change of state depends on the state. Why? Because the higher the water level, the greater the water pressure at the drain, and the faster the water will flow out. But as the water flows out, the pressure decreases, and so the flow rate also decreases. In order to make this into a real change

equation, we need a mathematical expression for how the flow rate depends on the water level X. As we just said, the greater the water level X, the greater the flow. Let's suppose that the relation is that they are proportional. What does that mean? To say that "Y is proportional to X" means there's a constant k such that $Y = kX$ (see Figure 1.25).

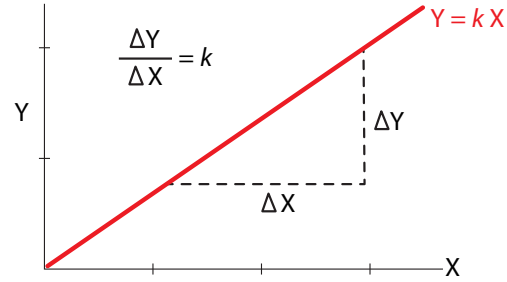

Figure 1.25: A proportional relationship; $k = \frac{\Delta Y}{\Delta X}$ is the slope of the red line.

Exercise 1.4.1 Write the following statements as equations.

a) A is proportional to B with a proportionality constant of 2.5.

b) X is proportional to Z with a proportionality constant of -3.7.

c) An animal's population density, P, is proportional to body size, B, with a proportionality constant of m.

In the case of the bathtub, the constant of proportionality is the width of the drainpipe. The wider the drainpipe, the faster the flow for a given pressure. So if $X' = -$outflow and outflow $= kX$, then we have the change equation

$$X' = -kX$$

In this equation, what are the units of k? Since X' is in gallons per minute and X is in gallons, the units of k must be $\frac{1}{\text{min}}$. **The units must always make sense in a change equation.** If the units don't match, we have to convert them so they do. For example, if X is in gallons and $k = 1\ \frac{\text{quart}}{\text{min}}$, then we have to convert k into $\frac{\text{gal}}{\text{min}}$ by multiplying $1\frac{\text{quart}}{\text{min}} \times 0.25\frac{\text{gal}}{\text{quart}} = 0.25\frac{\text{gal}}{\text{min}}$.

Of course, k is just a symbol. Let's assume it has the value $k = 0.2$ for this bathtub and drain. Then our change equation is

$$X' = -0.2X$$

Change Equations More Generally

We will now look at change equations more generally. The ingredients of such equations, which we will discuss now, are stocks and flows.

The Variables

The values of the quantities being modeled collectively make up the *state* of the system, and the quantities themselves are often referred to as *state variables*. State variables are *stocks*—

loosely speaking, accumulations of stuff. The amount of water in a bathtub, the amount of money in your bank account, the amount of energy in a battery, and the number of antelopes in a population are all stocks. In a system of change equations, the amount of a stock is a state variable.

Exercise 1.4.2 Give three more examples of stocks.

Change equations tell us how fast the state variables are changing and whether the change is positive or negative.

You should keep in mind that in most cases, rates of change of state variables are not themselves state variables. (We already saw one exception to this: in mechanics, velocities are state variables.) When identifying the state variables in a system, look first at stocks, not rates of change of stocks.

State variables vs. parameters in models In the bathtub model above, the state variable is X. Variables change their values as the system changes over time. But what about k? It is constant for a given model and doesn't change. It is called a *parameter* of the model. Parameters are, for right now, fixed numbers like 0.2. Later on, we can generalize this to parameters that change on their own with time (like an outflow tube getting narrower with time). In this text, we will always use capital letters for state variables, and lowercase letters for parameters, so as never to confuse them.[5]

As an example, consider a hot cup of coffee in a cooler room. Let's represent the state variable of this system by T = temperature of the coffee (in degrees Kelvin). Then Newton's law of cooling says that the change in temperature of the coffee is proportional to the difference between the temperature of the coffee and the temperature of the room. If we let the room's temperature be r (a parameter), then the difference between the coffee's temperature and the temperature of the room is $T - r$. So Newton's law of cooling gives us the change equation

$$T' = const \cdot (T - r)$$

where $const$ represents some proportionality constant, which will be another parameter in this model. Before we assign a name to this proportionality constant, we can say a little more about it based on intuition. Since the coffee cup is hotter than the room, we know that $T > r$, so $T - r$ is positive. But from basic intuition, what will happen to a hot cup of coffee in a cool room over time? The coffee will eventually cool off. In other words, in the language of our model, T will decrease. And what does this mean about T', the change in T? It means T' will be negative. If T' must be negative, but $T - r$ is positive, then in order for the change equation above to work, $const$ must be negative. (This intuition also works the other way around: if the coffee is actually an iced coffee in a warm room, then T would be less than r, so $T - r$ would be negative. But in this situation, the coffee would get warmer over time, meaning that T' would be positive. Once again, this means that $const$ must be negative.) Since $const$ must be a negative constant, we will replace it with $-k$, where k is a (positive) parameter. Thus, our final change equation is

$$T' = -k(T - r)$$

[5] This is just a convention in this textbook. Out "in the real world" (i.e., in most scientific fields), it is common for certain parameters to be capitalized, and sometimes lowercase letters will represent state variables. So remember that a more reliable way to distinguish state variables from parameters is this: if there is a change equation for something, that thing is a state variable. Otherwise, it's a parameter.

Inflows and Outflows

The most straightforward way to write change equations for a system of stocks and flows is to go through the system's stocks one by one and record the inflows and outflows of each stock. For a bathtub, the inflow is water flowing from the faucet, and the outflow is water flowing down the drain, as diagrammed in Figure 1.26. For a bank account, the inflows are deposits and interest, and the outflow is withdrawals. A nonrechargeable battery has no energy inflow, while the outflow is energy used to run the flashlight, radio, or other system the battery is powering. For an animal population, the inflows are birth and immigration (migration into the population), while the outflows are death and emigration (migration out of the population). These stocks and flows can be represented using the box-and-arrow diagrams in Figure 1.27.

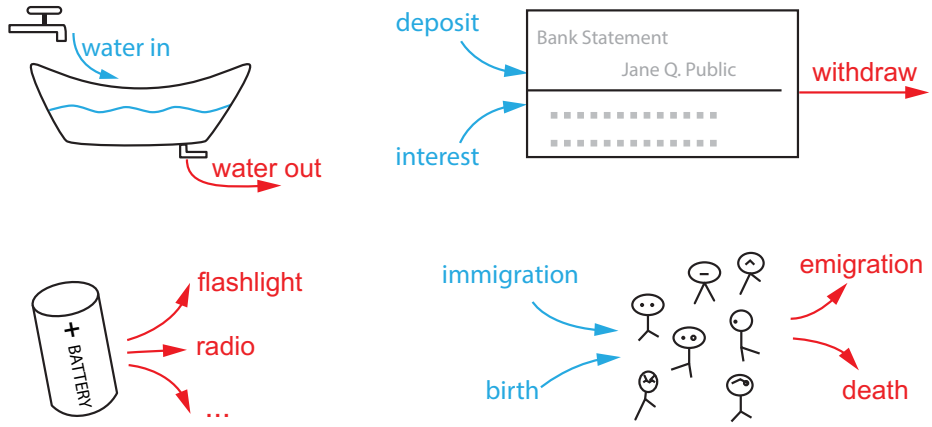

Figure 1.26: Systems with both inflows (blue) and outflows (red), except the battery (lower left) which has only outflows.

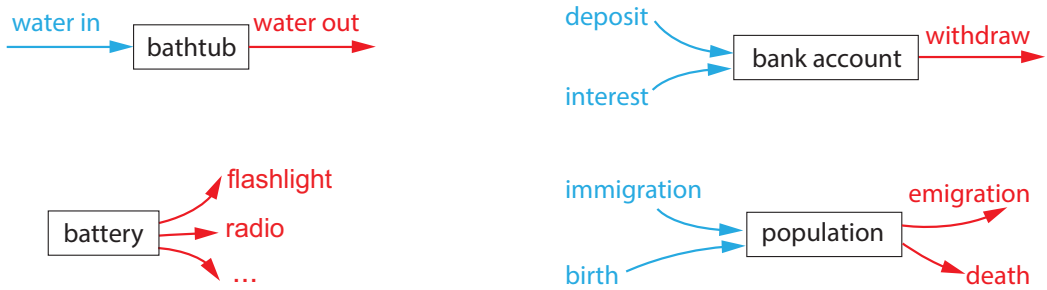

Figure 1.27: Schematic box-and-arrows diagrams of the systems in Figure 1.26.

Exercise 1.4.3 List all the inflows and outflows for each stock you came up with in the previous problem.

Exercise 1.4.4 Draw a box-and-arrow diagram for each of your stocks.

Once we know the inflows and outflows affecting a stock, we can write a word equation describing how the stock will change over time. These equations always have the following general form:

change in stock = input flows − output flows

The word equation for the bathtub example shown in Figure 1.26 is

$$\begin{aligned}\text{change in amount of water (per time)} &= \text{inflow rate} \\ &\quad - \text{outflow rate}\end{aligned}$$

The word equation for the population example shown in Figure 1.26 is

$$\begin{aligned}\text{change in population (per year)} &= \text{births per year} \\ &\quad + \text{immigrants per year} \\ &\quad - \text{deaths per year} \\ &\quad - \text{emigrants per year}\end{aligned}$$

Stocks don't always have both inflows and outflows. Sometimes, only one of these exists, as in the examples below.

A landfill receives inputs of trash, but none comes out. The word equation is

$$\text{change in amount of trash (per day)} = \text{trash dumped (per day)}$$

A box of tissues is used but never refilled. The word equation is

$$\text{change in number of tissues (per week)} = -\text{tissues used (per week)}$$

Notice that these word equations include only flows. There is **never** a separate term for the value of the stock, either at the current time or at the beginning of our observations. This is because when writing these equations, we consider only how the stock is changing, not its actual value. This seems counterintuitive at first but turns out to be a powerful way of modeling many kinds of systems.

Exercise 1.4.5 Write word equations for the bank account and battery examples in Figure 1.26.

Exercise 1.4.6 Write word equations for your three box-and-arrow diagrams.

From Words to Math

Once we have a word equation, we must then turn it into a change equation in mathematical form. To do this, we assign symbols to state variables and then write numbers or mathematical expressions for each flow.

We can denote the amount of water in a bathtub by the symbol W. Suppose the inflow is $2\ \frac{\text{gal}}{\text{min}}$ and the outflow is $1\ \frac{\text{gal}}{\text{min}}$. This gives the change equation $W' = 2 - 1$, or $W' = 1$.

Let T be the number of tissues in a box and suppose you use 7 tissues in an average week. This gives the change equation $T' = -7$.

Exercise 1.4.7 Why is there a minus sign in front of the 7 in the previous example?

Exercise 1.4.8 Call the amount of trash in a landfill L and suppose 1000 pounds of trash are added to the landfill daily. Write a change equation for the amount of trash.

Exercise 1.4.9 Suppose 100 births and 95 deaths occur in a population each year. Also, 3 individuals enter the population from outside and 2 leave. Write a change equation for the population size, P.

We will now examine several models from biology and other areas.

One-Variable Systems

A Simple Population Model

Think of an animal population, and let's say the state variable is X, the number of animals. What changes X? Well, one thing that changes X is animal births, and another is animal deaths. We can write a change equation

$$X' = \text{birth rate} - \text{death rate}$$

But how do we represent the birth and death rates? We are going to have to make some highly simplified assumptions. These assumptions are very strong and have huge consequences for the model.

All models make huge assumptions, and it is critical to be able to state what they are for a given model. The validity of a model depends strongly on its assumptions.

For example, for our first pass we are going to assume that animals don't die. (This might make sense if we are looking at a time frame much shorter than the typical lifetime). Then we have

$$X' = \text{birth rate}$$

How do we represent the birth rate? Let's make some huge assumptions: (1) there are no sexes; all animals are capable of giving birth, (2) an animal's ability to give birth is constant over its lifetime from birth to death (3) all animals have the same likelihood of giving birth.

Then for each animal, there is a single constant rate b at which that animal gives birth, let's say $b = 0.5$ babies per year (one baby every two years).[6] Then we say that the *per capita birth rate* is given by

$$\text{per capita birth rate} = b = 0.5$$

[6] These assumptions amount to saying that we can average varying birth rates over the whole population to produce a single number.

A "per capita" (literally "per head") birth rate is the rate of birth for a single animal. Its units are (animals per animal) per year, which is equal to $\frac{1}{\text{year}}$. This per animal rate must then be multiplied by the number of animals (X) to get the total birth rate. So we end up with

$$X' = bX$$

Let's consider another model, in which there is no birth, but animals die. So there is a death rate, and we will again make some highly simplified assumptions: (1) every animal has the same likelihood of dying, (2) the death rate does not depend on the number of animals, (3) the death rate does not vary with time. Then we can define the per capita death rate d, and write

$$X' = -dX$$

We could also combine birth and death in another model and write

$$X' = bX - dX$$

or

$$X' = (b - d)X = rX$$

where r is the net per capita growth rate.

To summarize, a model of a process is a **change equation**, in which the changes in a system depend on the current states. We write

$$X' = f(X)$$

Exercise 1.4.10 Write change equations for the following situations. You can use X or anything you prefer for your state variable.

a) A population has a per capita birth rate of 0.3.

b) A population has a per capita death rate of 0.4.

c) A population has a per capita birth rate of 0.25 and a per capita death rate of 0.15.

d) A population has a per capita birth rate of 0.1 and a per capita death rate of 0.2.

Exercise 1.4.11 In each of the cases in the previous exercise, is the population growing or shrinking?

A glimpse ahead How will we use the model to make predictions about behavior? We will start at a given state, which then gives the change (through the change equation), and then the change plus the present state will give us a new state, which will give us a new change of state, . . . etc. We will explore this process in more detail later in this chapter.

A Population Model with Crowding

As we said above, the model of rabbit population growth $X' = b \cdot X$ is pretty dumb if you take it too seriously at large values of X. You will see later that in the long run, this model predicts the existence of a ball of fifteen quintillion rabbits expanding outward at half the speed of light,

which is not exactly realistic. What is this model missing? Any effects of *crowding*, *such as competition for scarce* resources, which would limit growth. In a model with no limits to growth, growth will be unlimited. So we need a *crowding term*.

What would that look like? We will multiply the birth rate bX by a "crowding factor," which will be some number ≤ 1:

$$X' = bX \cdot \text{crowding factor}$$

To derive an expression for this crowding factor, let's suppose that the environment has a *carrying capacity* of k animals. Then $\frac{X}{k}$ would represent the fraction of the carrying capacity that is already being used by the present population X, which leaves $(1 - \frac{X}{k})$ as the fraction of resources that are currently unused and therefore available.

So our new change equation, including crowding, is

$$X' = bX(1 - \frac{X}{k})$$

There is another approach to same equation. Let's think about what the effect of crowding on X' is. It's certainly negative, so it has a minus sign. It certainly will get bigger as X (the population size) gets bigger. But we can be more specific than that. How often will an X bump into another X out at the lettuce patch? By analogy, think of a very large deck of cards, made up of many poker decks. In that large deck, what is the probability of drawing 2 aces? Probability theory tells us that the chance of drawing two aces from that large deck of cards is equal to the

$$(\text{probability of drawing an ace}) \times (\text{probability of drawing an ace})$$

or

$$(\text{probability of drawing an ace})^2$$

So the chance of two rabbits landing on one small lettuce patch, like the chance of drawing two aces, is proportional to the square of the number of rabbits (X^2). What is the constant of proportionality? Let's call it c, giving us

$$X' = bX - cX^2 \tag{1.1}$$

The crowding parameter c is therefore equal to $\frac{b}{k}$. This equation, which is called the *logistic equation*, is important in mathematical biology, and you will see it many times in this book.

Note that we can simplify equation Equation 1.1 by factoring out bX from the right-hand side:

$$X' = bX\left(1 - \frac{X}{k}\right)$$

Remember that positive change means *increase*, and negative change means *decrease*. It's interesting to consider when X' is positive and when it's negative. Since b (the per capita birth rate) is positive, and X (the population) is certainly always positive, the right-hand side of this equation will be positive when the term $(1 - \frac{X}{k})$ is positive and negative when it is negative. When X is smaller than k, the fraction $\frac{X}{k}$ will be smaller than 1, so $(1 - \frac{X}{k})$ will be positive. This means that when the population is smaller than the carrying capacity, X' will be positive, and therefore the population will increase.

Exercise 1.4.12 What happens when X is larger than k?

This basic reasoning helps to reassure us that our model behaves the way it should.

Two-Variable Systems

Romeo & Juliet

A wonderful set of examples, initially developed by Cornell mathematician Steve Strogatz, concerns the love dynamics of a couple we will call Romeo and Juliet. We will let R represent Romeo's feelings for Juliet, and J represent Juliet's feelings for Romeo. Positive values represent love and negative values represent hatred.

The state space for the Romeo–Juliet system is the 2-dimensional space (R, J).

What changes J and what changes R? That obviously depends on the details of their personality types and their relationship.[7] For example, let's assume that the changes in Juliet's love do not depend on her own feelings, but are purely a reflection of Romeo's love for her. If his love is positive, hers grows, and if he hates her, her love will decrease, possibly even into hate. Let's say the change in Juliet's love is proportional to Romeo's love. Let's assume, for this first model, that it's a linear proportionality, and that the proportionality constant is 1. So we have just said that $J' = R$.

Romeo, on the other hand, has issues. He also does not care about his own feelings and only reacts to Juliet, but in his case, the reaction is negative. If Juliet loves him, his love declines, and if she hates him, his love will increase. So $R' = -J$.

Our complete model is now given by the pair of change equations

$$J' = R$$
$$R' = -J$$

Exercise 1.4.13 Suppose that in addition to being turned off by Juliet's love, Romeo is turned off by his own love for her. Specifically, Romeo's love declines at a rate proportional to itself with proportionality constant k. Write a model for the Romeo–Juliet system that adds in this assumption.

Springs

Consider a basic example in mechanics: a simple mass–spring system (Figure 1.28). The mass is a cart that is attached to a spring, and rolls back and forth. What are the states of this system? Obviously, the position X of the mass (the cart) is one state variable, but are there any others? Yes! In physics, and in mechanics in particular, *velocity* is also a state variable, necessary to describe the state of the system. "The position of the mass is at $X = 2$" is one state, but "the velocity is $+3$" is also necessary to predict the system's future behavior. The mass being at $X = 2$ and heading to the left at 3 meters per second is in a different state from that in which the mass is at $X = 2$ and heading to the right at 5 meters per second. So in mechanics, state spaces tend to have two types of state variables: positions and velocities.

In this case, the state space of the mass-spring system is made up of all pairs (X, V) representing the position of the cart (X) and its velocity (V).

Now let's write the change equations

$$X' = f(X, V) \qquad \text{and} \qquad V' = g(X, V)$$

[7] See Strogatz's "Love affairs and differential equations," *Mathematics Magazine* (Strogatz 1988) or his book (Strogatz 2014) for some great examples. Generalizations to more complex psychologies and more than two people can be found in Sprott (2004) and Gragnani et al. (1997).

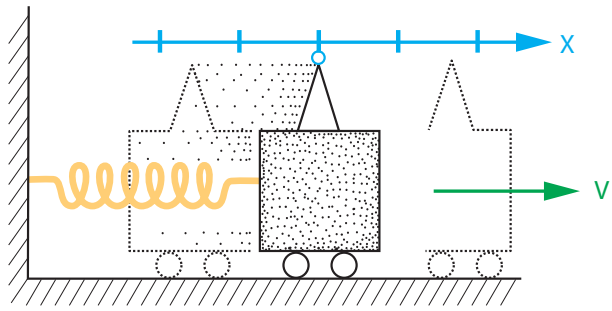

Figure 1.28: Mass–spring apparatus (adapted from Abraham & Shaw, *Dynamics, the Geometry of Behavior* Abraham and Shaw 1985). At any given time, the cart has a position X, given by the pointer, and a velocity V.

remembering that we are trying to see how each state variable changes depending on its own value and the values of the other state variables.

First, what changes position X? By definition, velocity *is* change in position, so $X' = V$.

How about V'? We have to recall a little physics here, specifically Newton's idea that the change in V, also called acceleration, is equal to the force applied to the object divided by the mass of the object. This is usually written as "$F = ma$," but that hides the fact that this is really a change equation. What it is really saying is

$$V' = \frac{F}{m}$$

(The mass m is a parameter in this model.)

But we're not done yet, because we still need to figure out what F is; F stands for the force acting on the object. What is this force? In this case, it is the force of the spring. And what is that? You may remember from high-school physics something called "Hooke's law," which says that the force of a spring is proportional to its extension and acts in the opposite direction: $F = -kX$. The proportionality constant k is called the *stiffness* of the spring (or simply the "spring constant"). Now, it turns out that Hooke's so-called law is false for most biological objects, such as muscles and tendons, and is true for metal springs only if they're stretched by small amounts. But let's assume that Hooke's "law" was really true and $F = -kX$.

Now we have a complete system of change equations

$$X' = V$$
$$V' = -\frac{k}{m}X$$

If we measure mass and spring stiffness in units for which $\frac{k}{m} = 1$, we get

$$X' = V$$
$$V' = -X$$

In other words, the simple spring has the same dynamics as our Romeo & Juliet example!

However, there is something not realistic about this spring model. Our model has not accounted for *friction*. In reality, there is always some friction, which changes the situation and changes the model.

How do we model friction? There are many different types of friction, ranging from air resistance to sliding friction to rolling friction to fluid viscosity, etc. We will make a very simple model of friction: that it is proportional to the velocity of the object. This is true, for example, for air resistance (think of riding a bicycle: the faster you go, the greater is the wind resistance) and for a viscous fluid.

In this case, we will model friction as a dashpot, a piston pushing though a fluid (Figure 1.29). So what is the force of friction? We will model it here as a simple negative force that is proportional to velocity.

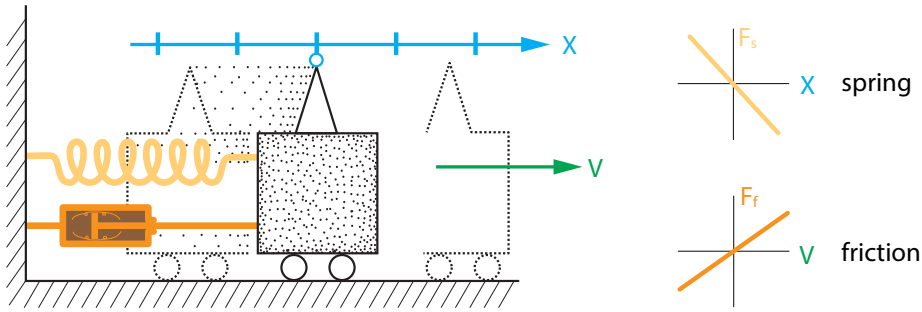

Figure 1.29: Left: The mass-spring apparatus now has a new element, signified by the dashpot (piston) attached to the cart. Right: The effect of the dashpot is to add a new force, friction. There are now two forces acting on the cart, the spring force, which is a linear function of position in this model, and the friction force, which is a linear function of velocity in this model.

Exercise 1.4.14 Write an expression for friction. You can make up parameters as necessary.

Exercise 1.4.15 Write the model for the spring with friction.

Sharks and Tuna

Let's develop a model of the shark–tuna system we have been talking about since the beginning of this book. We will call the number of sharks S and the number of tuna T, so a state vector for the system has the form (S, T). The model's state space is $\mathbb{R}_+ \times \mathbb{R}_+$, or the positive quadrant of $S - T$ space.

To start, let's write $S' = \ldots$ and $T' = \ldots$ and ask what changes S and what changes T.

What changes S? Sharks are born and sharks die. We will assume that sharks die at a constant per capita rate d. The shark birth rate, on the other hand, we are going to assume is proportional to the amount of food the sharks get. Let's call the proportionality constant m. It reflects the relative size of the tuna as food for the shark. If m is large, then one tuna makes a big difference to the per capita shark birth rate; if it is small, then the shark needs a lot of tuna to reproduce. So this term in the equation for S' is $m \cdot$ [available food]. But what determines the amount of available food? The tuna population! Every time a shark encounters a tuna, there is a certain probability that the shark is going to catch and eat the tuna. We will call that probability β (the Greek letter beta). This type of β parameter, which controls the frequency of successful (for the shark) shark–tuna encounters, is very common in all kinds of population modeling.

Another way to see why the shark–tuna encounter term involves the product ST is to see that the likelihood of a shark–tuna encounter depends on the probability of a shark finding a tuna in a given patch of ocean. Using the same reasoning as on page 30, we see that this is equal to the product of the probability of finding a shark times the probability of finding a tuna, or ST. So if β is the probability that a shark–tuna encounter results in the shark catching the tuna, we can write

$$S' = m\,\beta ST - dS$$

Similar reasoning tells us that the equation for the change in the tuna population is

$$T' = bT - \beta ST$$

Exercise 1.4.16 Describe this reasoning. Where does each term in the T' equation come from? What assumptions do we need to make to derive it?

These are called the *Lotka–Volterra predator–prey equations* for their two (independent) discoverers, Alfred Lotka and Vito Volterra. They were developed in the early twentieth century.

Clearly, a lot depends on d, b, m, and β. Soon, we will develop the tools study this. For now, we are going to set all these parameters to 1 (with appropriate units so that addition and subtraction make sense). We thereby lose all quantitative validity but gain a qualitative view of the model, which is now simply

$$\begin{aligned} S' &= ST - S \\ T' &= -ST + T \end{aligned} \tag{1.2}$$

This model produces oscillatory behavior of the shark and tuna populations (Figure 1.2).

Chemistry

In chemistry, we learn that chemical reactions are written as

$$\mathrm{A} + \mathrm{B} \xrightarrow{\mathrm{k}} \mathrm{C}$$

meaning "A plus B yields C with rate constant k." Since the amounts of A, B, and C are changing, it should be possible to write change equations for them. However, the arrow ("goes to") is not a mathematical symbol, so how do we translate this into math?

Let's start with a simpler example,

$$\mathrm{A} \xrightarrow{\mathrm{k}} \mathrm{B}$$

The state variables in this system are [A], the concentration of A, and [B], the concentration of B. To avoid having to write all those square brackets, we will write "A" to mean "[A]."

Let's start by looking at chemical A. To find A', we ask, "what makes A go up, and what makes A go down?" In this particular reaction, nothing makes A go up; A can only go down, and it goes down when molecules of A turn into molecules of B, which happens at a rate k. "A turns into B at a rate k" means that in one unit of time, the fraction of A molecules that turn into B molecules is k. In other words, the situation is exactly analogous to the "per capita death rate" in our population models: k is a "per molecule death rate" (or rather, a per molecule rate of A turning into B). Therefore, k must be multiplied by the number of molecules of A to determine

the total change. But the number of molecules is just A. (We said earlier that the variable A is the concentration of A molecules, but of course if the volumes are held constant, then changes in [A] are just changes in the number of molecules.) Therefore, we have the change equation

$$A' = -kA$$

Exercise 1.4.17 Use similar reasoning to write the equation for B'.

Now let's look at the slightly more complicated reaction

$$2\mathrm{A} \xrightarrow{\mathrm{k}} \mathrm{B} \qquad or \qquad \mathrm{A} + \mathrm{A} \xrightarrow{\mathrm{k}} \mathrm{B}$$

Again, the state variables will be A and B. Also, as before, nothing makes A go up, and A goes down when A turns into B. In this reaction, however, two molecules of compound A must collide in order to form a molecule of B. How often will that happen? The *law of mass action* from chemistry tells us that chemicals participate in chemical reactions in proportion to their concentrations. This means that the frequency with which a molecule of A will bump into another molecule of A is proportional to the square of its concentration, or A^2. Thus

$$A' = -2kA^2$$

This reasoning might sound familiar. It's exactly the same logic as that of the logistic population growth model, where we said that the frequency with which rabbits bump into other rabbits is proportional to the square of the rabbit population. (The "2" on the right-hand side is there because two molecules of A combine to form each molecule of B, so each successful collision removes two molecules of A.) We will soon see another analogy between models in chemistry and ecology.

Exercise 1.4.18 What is the equation for B' in this case?

Now we can return to the reaction that began this section,

$$\mathrm{A} + \mathrm{B} \xrightarrow{\mathrm{k}} \mathrm{C}$$

The state variables in this system are A, B, and C. Let's start by looking at chemical A. In this reaction, nothing makes A go up. What makes A go down is A combining with B to make C. So, how often will an A molecule bump into a B molecule? The frequency of collision is proportional to the number of A molecules and the number of B molecules, and therefore proportional to their product. The rate constant is the constant of proportionality, so we can write

$$A' = -kAB$$

This is exactly the same logic as that of the shark–tuna model, where the frequency of shark–tuna encounters is proportional to the product of S and T, where S is the shark population, and T is the tuna population. Now, if we divided these populations by the volume of the ocean patch we are modeling, they would become the shark and tuna concentrations (i.e., population densities). In this example, the proportionality constant k takes the place of what we previously called β: the probability that a shark–tuna encounter results in the death of the tuna. In the chemical reaction case, k is called the *rate constant*, and it has basically the same interpretation. In fact, the shark–tuna system can be rewritten as a set of chemical reactions (with D representing dead sharks):

$$\mathrm{S} + \mathrm{T} \xrightarrow{\beta} 2\mathrm{S}$$

$$S \xrightarrow{d} D$$

$$T \xrightarrow{b} 2T$$

Exercise 1.4.19 Following this analogy, finish writing the change equations for the reaction $A + B \xrightarrow{k} C$.

Let's go back to chemical reactions and look at the reversible case, in which

$$A + B \underset{k_b}{\overset{k_f}{\rightleftharpoons}} C$$

Now something does make A go up, namely, the back-reaction of C dissociating into A and B. By the same logic as before, we get the change equation

$$A' = k_b C - k_f AB$$

Exercise 1.4.20 Write the change equations for B and C.

Note that the equation for A' has an interesting implication. We talk about chemical "equilibrium." As we will see in Chapter 3, at equilibrium, by definition, there is no net change in the concentrations. That means that $A' = B' = C' = 0$. Chapter 3 will exploit the structure of equilibria, but here is an immediate application. In chemistry, we are taught that at equilibrium, the final concentrations in a chemical reaction stand in a certain ratio. This ratio can be derived from the conditions for $A' = 0$.

So if $A' = 0$, then $0 = k_b C - k_f AB$. Therefore, at equilibrium,

$$\frac{C}{AB} = \frac{k_f}{k_b}$$

A Model of HIV Infection within an Individual Person

In 1981, patients with strange infections and cancers started showing up at UCLA Medical Center. Soon, similar cases were identified elsewhere, and the disease was given the name AIDS, for acquired immune deficiency syndrome. In 1983, the virus that caused AIDS was identified and, a few years later, named human immunodeficiency virus, or HIV.

HIV infects a particular type of white blood cell called a $CD4^+$ T lymphocyte. When a cell is infected, there are two possibilities. If the cell is actively infected, it starts budding off viruses and dies within a few days. If it's latently infected, viral genes are incorporated into the cell's genome. The cell remains healthy, but the infection can become active at some point in the future.

When a person is first infected with HIV, the amount of virus in their blood goes up to a very high level. This lasts for a few months, and then their virus count goes down to a fairly low level. Initially, doctors thought this happened because of an immune response to the virus. However, in 1996, the mathematical biologist Andrew Phillips published a paper asking whether this spike-and-decline pattern was possible even without an immune response (Phillips 1996). To do this, he used a model of HIV infection developed a few years earlier by Angela McLean and her colleagues (McLean et al. 1991).

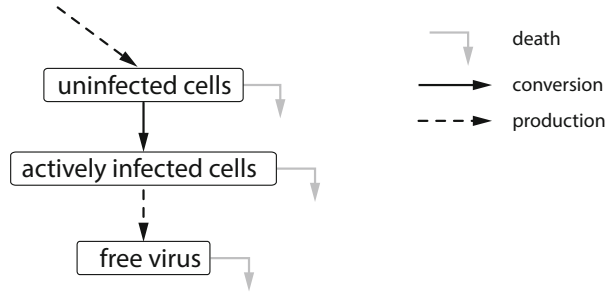

Figure 1.30: A simplified version of the McLean–Phillips model that leaves out latently infected cells.

We're going to build a slightly simplified version of the McLean–Phillips mathematical model of what happens in the bloodstream of a person infected with HIV. This model will consist of three equations showing how the sizes of the different populations in Figure 1.30 change over time. The model variables are the amounts of virus, uninfected cells, and infected cells. We'll call them V, R, and E, respectively. We now need to write a change equation for each one of these variables.

Viruses

Viruses are produced by infected cells. Once produced, they can die (become noninfectious) or infect new cells. However, such a small fraction of virus particles infects new cells that we're going to assume that this doesn't affect the amount of virus in the bloodstream. We can write a word equation for the change in amount of virus in an infected person's bloodstream:

$$\begin{aligned}\text{change in amount of virus (per day)} &= \text{viruses produced per day} \\ &\quad - \text{viruses dying per day}\end{aligned}$$

Let's look at each term on the right-hand side of this word equation. First, we have the number of viruses produced each day. On average, each infected cell produces 100 virus particles every day. Another way to put this is that the *per capita production rate* of viruses by infected cells is 100.

If the per capita virus production rate is 100 per day, then the expression for how many viruses are produced each day is $100E$. (Write this expression and those derived later on the appropriate arrows in Figure 1.30.) Next, the per capita virus death rate is 2 per day, meaning that an average virus lives only half a day. The total number of virus deaths per day is then $2V$. Therefore, the full equation for the rate of change of the virus population is

$$\underbrace{V'}_{\text{change in amount of virus (per day)}} = \underbrace{100E}_{\text{viruses produced per day}} - \underbrace{2V}_{\text{viruses dying per day}} \tag{1.3}$$

Uninfected Cells

Uninfected cells are produced by the body, die natural deaths, and can become infected by HIV. The word equation is

$$\begin{aligned}\text{change in uninfected cells (per day)} &= \text{cells produced per day} \\ &\quad - \text{cells dying per day} \\ &\quad - \text{cells infected per day}\end{aligned}$$

Let's once again translate this word equation into math. First, the rate of production of uninfected cells is 0.272 per day. (That may seem unrealistically small, but to keep the numbers manageable, we're simulating what happens in $1\,\text{mm}^3$ of blood. The total rate of $CD4^+$ T lymphocyte production for the whole body is correspondingly larger.) The per capita death rate for uninfected cells is 0.00136, so the total death rate is $0.00136R$.

Now we need to consider infection. For an uninfected cell to be infected, it must encounter a virus particle in the bloodstream. As with the shark–tuna model and the chemical reaction rate models we developed, the chances of this happening are directly proportional to the *product* of the amount of virus and the number of uninfected cells in the bloodstream. Translating this into a math expression, the infection rate is βRV, where β is the proportionality constant. It turns out that β is about 0.00027. Thus, the full equation for the rate of change of the amount of uninfected cells is

$$\underbrace{R'}_{\text{change in uninfected cells (per day)}} = \underbrace{0.272}_{\substack{\text{cells produced}\\ \text{per day}}} - \underbrace{0.00136R}_{\substack{\text{cells dying}\\ \text{per day}}} - \underbrace{0.00027RV}_{\substack{\text{cells infected}\\ \text{per day}}} \tag{1.4}$$

Infected Cells

Finally, we'll write the equation for infected cells. These cells arise from infection of uninfected cells, and all of them eventually die from the infection. So the word equation is

$$\begin{aligned}\text{change in infected cells (per day)} = {}& \text{cells infected per day}\\ & - \text{infected cells dying per day}\end{aligned}$$

To turn this word equation into math, the per capita mortality rate for infected cells is 0.33, which means that about one-third of the cells die every day. Thus the total mortality rate for these cells is $0.33E$. And finally, we already know the rate at which uninfected cells become infected. From the discussion of uninfected cells above, this rate is $0.00027RV$. Therefore, the last change equation in the model is

$$\underbrace{E'}_{\text{change in infected cells (per day)}} = \underbrace{0.00027RV}_{\text{cells infected per day}} - \underbrace{0.33E}_{\text{infected cells dying per day}} \tag{1.5}$$

Our model of an HIV infection is now complete. To summarize, the three equations Equations 1.3, 1.4, and 1.5 form our complete system of change equations:

$$\begin{aligned}V' &= 100E - 2V\\ R' &= 0.272 - 0.00136R - 0.00027RV\\ E' &= 0.00027RV - 0.33E\end{aligned}$$

System Behavior

Spike and decline The reason we went to the trouble of writing these equations is that they can tell us what the consequences of our biological assumptions are. We wanted to know whether a decline in HIV levels could occur without an immune response. Since there's nothing about an immune response in the assumptions that led to our equations, if we observe a spike and decline, we'll know that an immune response isn't necessary for this.

Using differential equations to look at behavior over time in this way is called *simulation*. The basic idea is that if we know the initial values of all the state variables, and we know how the system changes at every point in time (i.e., the change equations), we can figure out how it will

behave if our assumptions are correct. You'll learn how to set up such simulations yourself a bit later in this chapter, but for now, you can use a prebuilt one.

Exercise 1.4.21 Run the three-compartment HIV simulation on the course website. Describe what you see happen.

This result tells us that you don't need an immune response to get a sharp drop in HIV levels after infection. It is exactly analogous to the drop in shark population that is seen when the tuna have been depleted. After a while, there just aren't enough susceptible cells to infect, which causes the decline. That doesn't mean that there isn't an immune response, but the drop in virus levels should not be taken as evidence of one.

Long-term behavior and model limitations If a person infected with HIV isn't treated, they eventually go on to develop AIDS, and their T cell levels fall far below what we're seeing in the simulation output.

Exercise 1.4.22 Does our model reproduce the drop in T cell levels seen when an HIV patient develops AIDS? It typically takes about ten years to develop AIDS. Run the simulation for 3650 days and describe the long-term behavior of the state variables.

We see that in this model, people infected with HIV don't get AIDS. But in real life, they obviously do. This tells us something very important: the progression from HIV infection to full-blown AIDS involves biological processes that this model doesn't describe. Just having cells getting infected and dying isn't enough; there needs to be more going on. What that "more" is, is a biological question. This is where modeling can interact with clinical and laboratory research in interesting ways.

There are models now that incorporate biology that our model doesn't, and those models demonstrate a progression to AIDS. But the simple model is still good for many purposes, like explaining the spike and decline in virus levels after infection. It can also be used to test other ideas, as you will see in the exercises.

Epidemiology

In the study of disease transmission in populations, modeling has become an important practical tool. Early in this chapter we saw the results of the model that the CDC used to predict the course of the Ebola epidemic (Figure 1.8 on page 6). The type of model that the CDC used is called a "susceptible–infected" model, or sometimes an "SIR" model (where R stands for "recovered").

One of the early models used to study the epidemiology of HIV transmission was the model of Anderson and May (Anderson et al. 1992). We present here a slightly simplified version of their model (Figure 1.31).

We will assume three populations:

S Susceptible individuals, that is, people who are HIV negative
I Infected-but-not-yet-symptomatic individuals, who are HIV positive, and
A People with the symptoms of AIDS.

We assume a fixed population of 10,000 people. Assuming that the average life span is 75 years, we would expect 1/75 of the population to die each year, giving a person's probability of

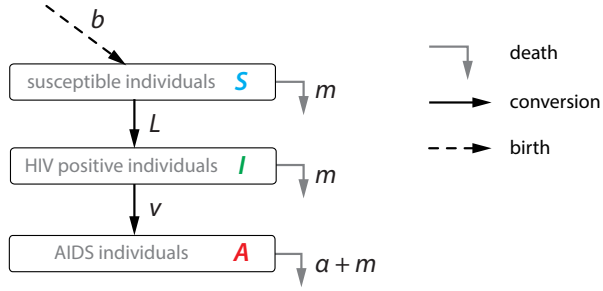

Figure 1.31: Schematic box-and-arrow diagram of a simplified version of the Anderson-May SI model for HIV transmission in a population.

dying is 1 in 75 years, giving a per capita death rate of $m = 1/(75 \text{ years})$. (Note that we are assuming that a person's probability of death is uniform across all ages, which is a limitation of this model. More advanced "age-structured" models use age-specific death rates.)

To compensate for these deaths, we also assume that $b = 133.3$ people are injected into the population each year, exactly making up for the natural death rate of $1/(75 \text{ years})$.

Exercise 1.4.23 Why 133.3? Where does this number come from?

The critical dynamical term is the susceptible-meets-infected term, which will have the form $S \times I$, just as it was with the sharks and tuna. Our underlying model here is a particularly simple one: we assume random encounters between members of S and I indiscriminately. In other words, we assume that neither party knows that an I is an I, that is, no one knows who is HIV+. Therefore, the probability that we encounter an I is $I/(S+I)$. We also assume that each person has, on average, c partners/year.

Of course, not every encounter between an S and an I results in the infection of the S. Just as in the shark–tuna model, there is a certain probability, which we call β, that the encounter will end in an infection. The parameter β is obviously extremely important: it is the parameter we can manipulate with safe sex practices and medications that reduce viral load and make infected people less likely to infect others. Let's begin by assuming that the probability of transmission of HIV with each encounter is a gloomy $\beta = 0.5$, or 50%.

Consequently, the overall per capita rate at which an S converts into an I is

$$L = c\beta \frac{I}{S+I}$$

We also need to reflect the fact that AIDS patients die more quickly than the average death rate. We assume an average AIDS-specific death rate of $\alpha = 1/(1 \text{ year})$. (This rate was more typical of the early days of the AIDS epidemic than it is today.) There is also a rate of conversion of I into A, that is, a rate of HIV+ people turning symptomatic, which we assume to take 8 years, giving a rate of conversion $I \to A$ of $1/(8 \text{ years})$ or 0.125/year.

The differential equations are therefore

$$\begin{aligned} S' &= b - (m+L)S \\ I' &= LS - (m+v)I \\ A' &= vI - (m+\alpha)A \end{aligned}$$

where

$$b = 133.33 \qquad m = 1/75 \qquad v = 0.125 \; L = c\beta\frac{I}{S+I}$$
$$\alpha = 1 \qquad c = 2 \qquad \beta = 0.5$$
$$\text{initial conditions : } S(0) = 9995 \; I(0) = 5 \quad A(0) = 0$$

With these parameters, the model predicts that the populations will go to equilibrium values at approximately

$$S = 152 \qquad I = 949 \qquad A = 117$$

This is a very gloomy outcome: almost 9000 out of our original 10,000 have died after 20 years, with most of that coming in the first 10 years (Figure 1.32).

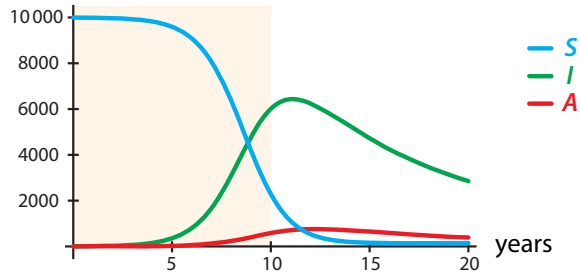

Figure 1.32: Time series output of the Anderson-May HIV model, assuming a high value for β. Note the outcomes.

But if we can change parameters, we can change outcomes. If we can lower β, for example by safe sex practices, to 0.05, the epidemic will die out. With this new value of β, a new equilibrium is reached, at approximately

$$S = 9994 \qquad I = 0 \qquad A = 0$$

indicating that we have prevented the virus from spreading (Figure 1.33).

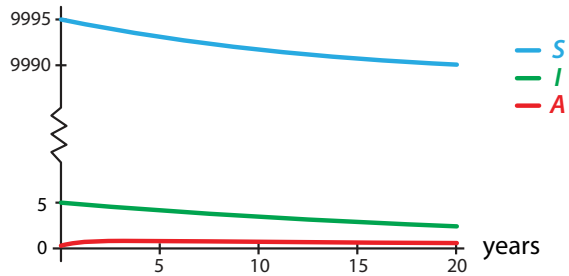

Figure 1.33: By lowering β, we can change the course of the epidemic.

Differential Equations

We have been talking about "change equations." The official, fancy, name for these is *differential equations*. The shark–tuna model, the Romeo and Juliet model, the spring model, the HIV model, etc., are all examples of differential equations.

Further Exercises 1.4

1. Translate the following verbal statements into differential equations. Use diagrams as necessary.
   a) The daily rate of change of B is 100.
   b) The yearly rate of change of P is -5.
   c) The monthly rate of change of H is 0.02 times H.
   d) The yearly per capita rate of change of G is -0.07.
   e) The weekly rate of change of L is the sum of an inflow, 0.05, and an outflow, 0.09.
   f) The daily rate of change of K is the sum of births, with per capita birth rate 3, and deaths, with per capita death rate 2.
   g) The daily rate of change of P is the sum of births, with per capita birth rate 2.5, deaths, with per capita death rate 1.3, immigration, with rate 10, and emigration, with per capita rate 0.6.

2. Here is another way to derive the logistic model you first saw on page 31. Consider an area of grassland that has enough resources to support a buffalo population of size k (k is called the grassland's *carrying capacity* for buffalo). The key assumption that this model makes is the following: *the per capita rate of change of the population is proportional to the fraction of resources available.*
   a) The fact that the carrying capacity is k means that when there are exactly k individuals in the population, they are utilizing 100 % of the resources. In that case, what fraction of the resources are used by one individual?
   b) If the current population size is X individuals, what fraction of the resources are they using collectively?
   c) In the same situation, what fraction of the resources are *not* being used?
   d) The per capita rate of change of X can be written as $\frac{X'}{X}$. The key assumption mentioned above says that this quantity is proportional to the expression you wrote down in part (c). Call the proportionality constant r, and write an equation for the per capita growth rate.
   e) Convert this equation into a differential equation (change equation) for the population's size. (Just make X' stand alone on the left-hand side of the equation.)

3. On a hot day, students are lining up to buy ice cream. Let L be the number of people in line. Write a differential equation for L using the following assumptions.

- Students join the line at a rate proportional to the number of people already in line, with a proportionality constant of 0.1.
- Students get ice cream and leave the line at a constant rate of 0.4 per minute.
- Students get tired of standing in line and leave at a per capita rate proportional to the number of people in line, with a proportionality constant of 0.02.

4. Collagen is a key protein in connective tissues. One of the steps in collagen formation involves the combination of three molecules of a collagen precursor called propeptide. This occurs with rate constant k. The rate of formation of propeptide is a constant, f. The propeptide also degrades with per molecule degradation rate d. Write a differential equation for the concentration of propeptide, P.

5. Mitochondria are organelles that provide energy for human and other eukaryotic cells. Mitochondria can divide like bacteria and fuse with each other. Use the following assumptions to write a differential equation for M, the number of mitochondria in a cell.
   - There is an optimal mitochondrial population, m. The rate at which mitochondria reproduce is proportional to the *difference* between the current population and the optimal population, with proportionality constant r.
   - When two mitochondria are close to each other, they may fuse together. This occurs with probability f.
   - Mitochondria die at a constant per capita rate d.

6. Spotted owls (W) prey almost exclusively on red-backed voles (V). Use the following assumptions to write a differential equation model of this system.
   - The vole population has a per capita birth rate of 0.1 and a per capita death rate of 0.025.
   - The rate at which an individual owl eats voles is proportional to the vole population with a proportionality constant of 0.01.
   - The owl birth rate is proportional to the amount of food they consume, with a proportionality constant of 0.05.
   - Owls have a constant per capita death rate of 0.1.

7. Przewalski's horse, a wild horse that inhabits central Asia, is the only horse species never to have been domesticated. In the wild, these horses are preyed upon by wolves. Write a model of the populations of Przewalski's horses (P) and wolves (W) based on the following assumptions.
   - The horse per capita birth rate is 0.15.
   - The horse per capita death rate is proportional to the population size, with proportionality constant 0.01.
   - Wolves prey on many species other than horses, so their per capita birth rate can be modeled as a constant, 0.1.
   - Wolves have a constant per capita death rate of 0.05.
   - A horse's probability of being eaten by a wolf is proportional to the number of wolves, with a proportionality constant of 0.02.

8. Kelp (K), sea urchins (U), and sea otters (S) form a food chain off the coast of northern California. Use the following assumptions to write a differential equation model of the food chain.
   - Kelp grows at a per biomass (like per capita) rate of 0.02.
   - Due to shading, kelp dies at a per biomass rate proportional to the amount of kelp with a proportionality constant of 0.01.
   - Sea urchins eat kelp. A single sea urchin consumes kelp at a rate of 0.05 per month.
   - The sea urchin birth rate is proportional to the amount of kelp the urchins consume, with a proportionality constant of 0.2.
   - Sea urchins die of natural causes at a per capita rate of 0.01.
   - The rate at which a single sea otter eats sea urchins is proportional to the sea urchin population with a proportionality constant of 0.03.
   - The sea otter birth rate is proportional to the amount of sea urchins the otters consume, with a proportionality constant of 0.01.
   - Sea otters die at a per capita rate of 0.001.

9. The pier in Santa Monica, CA, is a popular destination for both tourists and locals. Visitors ride the Ferris wheel (F), eat ice cream (C), or just walk around on the pier (W). Write a dynamical model for the numbers of people engaged in these activities given the following assumptions. *(Hint: Start by drawing a diagram of this system and labeling the stocks and flows.)*
   - People entering the pier always start by just walking around.
   - E people enter the pier each minute.
   - Visitors leave at a constant per capita rate d. They can leave only when they are walking around.
   - Due to fear of nausea, people do not go directly from eating ice cream to riding the Ferris wheel.
   - Visitors prefer to go on the Ferris wheel with friends. Thus, the probability that any one individual will go on the Ferris wheel is proportional to the number of people walking around, with proportionality constant b.
   - Riders leave the Ferris wheel at per capita rate n.
   - When visitors leave the Ferris wheel, a fraction z of them go directly to eating ice cream. The others walk around.
   - Visitors who are walking around prefer to avoid long lines for ice cream. Thus, the per capita rate at which they get ice cream is proportional to the inverse of the number of people already doing so, with proportionality constant m.
   - People who are eating ice cream stop doing so at a constant per capita rate k.

10. A simple model of infectious disease spread is

$$\begin{aligned} U' &= vW - mU - pU \\ V' &= qV - mV - rVW \\ W' &= rVW - mW - vW \end{aligned}$$

where v, m, p, q, and r are positive constants. The variables U, V, and W stand for susceptible population, infected (but not symptomatic) population, and symptomatic population (but not in that order).

a) Which variable represents which population? Justify your answer.

b) How would you model

1. a "safe-contact" program that reduced the probability of infection per encounter?
2. a drug that slowed the progression from infection to the appearance of symptoms (as AZT did for HIV/AIDS)?
3. a drug that cures the disease, in the sense that it makes infected people fully recovered but not immune?

11. In this textbook, we will use capital letters for state variables and lowercase letters for parameters, but many models in the scientific literature don't follow this convention. Both state variables and parameters can be written as either uppercase or lowercase letters. In the examples below, identify the state variables and parameters. *(Hint: Think about what state variables do that parameters don't, or see the footnote on page 26.)*

a) $g' = 0.2g$

b) $a' = 0.35ab$, $b' = -2b$

c) $X' = aX + RW$, $W' = RX$

d) $c' = Qcd - Rd$, $d' = Pd - Rc$

12. The HIV model studied in this section ignored the fact that some cells infected with HIV become latently, not actively, infected. If a cell is latently infected, viral genes are incorporated into its genome. The cell remains healthy for a time, but the infection can become active at some point in the future. (The original McLean–Phillips model included latent infection.) In this exercise, you will extend the model developed in the text to include latency.

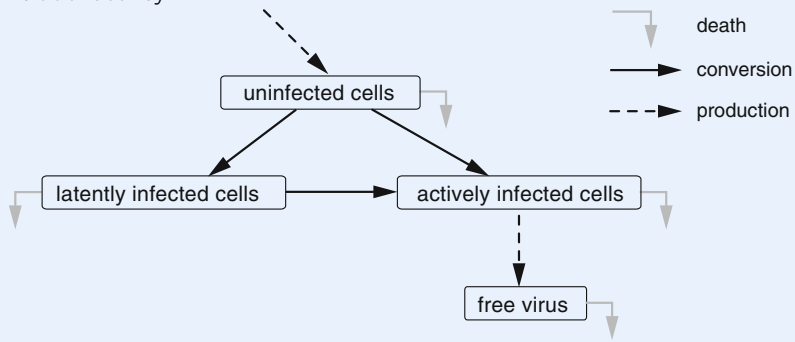

a) Write a word equation for each box in the above graph. Distinguish between latent and active infection.

b) Assume that 90 % of infected cells become actively infected and 10 % become latently infected. The per capita "activation rate," the rate at which latently infected cells become actively infected, is 0.036. Also, latently infected cells have the same

per capita death rate as uninfected ones. Use this information and the text to label the above graph with all state variables and flow rates.

c) Translate the word equations you wrote in part (a) into differential equations.

d) A simulation of this model is available on the course website. Run it and describe what happens. How much difference does the inclusion of latency make to the dynamics you observe?

13. The current treatment for HIV is a multidrug regimen that can reduce patients' virus levels to the point of undetectability. These drugs work by making it much harder for viruses to infect cells. We can simulate this with our model.

a) First, we'll make a change in notation. So far, all the parameters in the model have been numbers, but once we want to manipulate them, it becomes easier to change some of them to symbols. Actually, in most models in mathematical biology, all the parameters are symbols. The parameter we want to change is the rate at which viruses infect cells. This is the parameter that we previously called β. Rewrite your equations and the diagram, changing 0.00027 to β.

b) What happens to a person's long-term T cell levels and viral load when we manipulate β? We're no longer interested in the spike behavior, so change the model's initial conditions to the values at which the simulation settled. Try several values of β and observe the resulting values of the state variables.

c) The drugs in current use can keep people with HIV alive for very long periods, but they aren't a cure. If the person goes off the drugs, their virus levels go back up. This is mainly because of the latently infected cells, which can persist for years and aren't affected by current treatments. They function like time-release HIV. What do you think would happen to a patient who received both the current drugs, which mostly keep HIV from infecting new cells, and a new treatment that caused latently infected cells to become actively infected more quickly? (Please answer this before simulating the situation.)

d) In the simulation you are working with, the activation rate is represented by the parameter α (the Greek letter alpha). Restore β to its original value and manipulate α. Describe the effect of your manipulations on the state variables' long-term values.

e) What happens if we raise α and lower β at the same time, simulating the effect of combining conventional therapy with a new one that raises the activation rate? Try several combinations of values and describe what happens. *(Hint: It's often useful to push the boundaries of a model, trying very high or low values.)*

While a cure for HIV is still a long way off, this approach to treatment, termed immune activation therapy, is currently an active research area.

1.5 Seeing Change Geometrically

The Notion of Tangent Space

We have now learned to describe the causes of change. The change equation $X' = f(X)$ says, "if you are in state X, then you are changing at rate $f(X)$." Through the function f, the model gives us, for every possible value of the state variable X, the change X' at that state value. For this reason, we will think of X' as giving a *change instruction*.

For example, the bathtub model change equation $X' = -0.2X$ says that if you are at $X = 20$ gallons, then your change is -4. But -4 what? The answer is -4 gallons per hour. So the change equation is to be read as saying that if you are at $X = 20$ gallons, then change by a rate of -4 gallons per hour.

Just as we defined the set of all possible values of X to be the *state space* of the model, we now want to think about the set of all possible values of X', that is, all possible change instructions. For reasons that will become clear soon, we will call this the *tangent space* of the model.

Recall (from Section 1.3) that the state space is the set of all possible values of the state variable X. What does the set of all possible change instructions (i.e., values of X') look like? Notice that it can't be the same as the state space, because the units of X and of X' are different: X is in gallons (or animals, or glucose concentration, or . . .), whereas X' is in gallons per hour (or animals per year, or glucose per hour, or . . .). Also, their respective values may be different. If G = glucose concentration, then G must be greater than or equal to 0; negative glucose concentrations do not make sense. But glucose *changes* can be negative! Recall that a negative change means that G decreases (and a positive change means that G increases).

Tangent Space: Geometric Version

Suppose our model is of a single animal population X. Then *geometrically*, we think of the state space of X as the positive half (right half) of the real number line, called $\mathbb{R}_+$ (Figure 1.17 middle row):

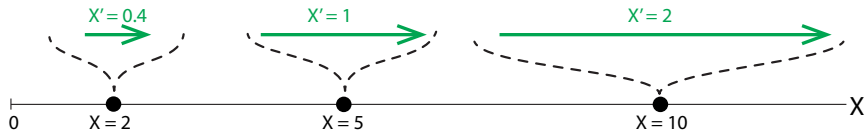

Figure 1.34: Three change vectors for the vector field $X' = 0.2X$.

What's the tangent space for this model? It's the whole of $\mathbb{R}$, positive and negative, because changes can be positive or negative.

Here is a device we are going to use heavily throughout this book. Remember that change is movement in state space. Therefore, it makes sense to think of the change instruction at a point as an arrow that points in the direction of the change and whose size indicates the magnitude of the change. And recall from Section 1.3 that arrows can be described by *vectors*. We will therefore refer to the change instruction that we get from X' as a *change vector*. In one dimension, the vector is pointing in the positive direction (to the right) if the change is positive (X is increasing) and in the negative direction (to the left) if X is decreasing. The length of the vector will represent the magnitude of the change. So, for example, in the population model

$X' = 0.2X$, at the point $X = 10$ animals, the change vector is pointing to the right and has length 2 animals per year (Figure 1.34).

As soon as we draw the change vectors in this way, we can grasp how the system is going to change. No matter where we start it, the change arrows point to the right, and they keep getting bigger and bigger. We can immediately say that the number of animals will grow without bound, and the rate of growth will get larger and larger as the numbers get bigger and bigger.

The next step is to think of the change vector as being superimposed on the state value to which it corresponds, as in Figure 1.35. This is a little bit of a fiction, but it is a very useful one. It's a fiction because the change vectors aren't really *in* state space. Rather, they are *assigned to* points in state space by the change equation.

Figure 1.35: Writing the change arrows (green) directly on the state space (black) is a useful fiction, very helpful to visualize how state points move through state space.

Vector Fields

We now have the key idea of this course: the model, which is a differential equation, gives us a change vector (a value of X', in the tangent space) corresponding to every state point (value of X) in the state space. In other words, the change equation is a *function* from the state space to the tangent space, which assigns a change vector to each state X. This view of the change equation as a function is so important that we give it a special name. We call it a *vector field*. What we have said is that a vector field is a function

$$\text{vector field:} \quad \text{state space} \to \text{tangent space}$$

We will use a standard color convention here: the state space is in black, and the change vectors are in green.

A vector field like the one shown in Figure 1.36 is very suggestive of movement, and indeed, it tells us how the state point moves through state space. Imagine the point being carried along by the vectors. Since there is a vector at every point, you can think of a crowd of people passing a beach ball overhead, with each person giving the ball a nudge in a particular direction.

Notice that there is a little problem with our picture of assigning the change arrows to the points. Since there is a different change arrow at *every* point, it's going to get awfully crowded in there, with change vectors looking like they are overlapping each other. But they're not! Keep in mind that the change vectors don't actually live in X space; they come from X' space, the tangent space. But when we draw a graphical representation of the vector field, we can't possibly draw every change vector, so we just draw some of them. Also, we superimpose them on the state space, even though they actually belong to the tangent space.

Let's look at the vector field for the logistic population growth model in Equation (1.1) on page 31:

$$X' = rX(1 - \frac{X}{k})$$

As we saw, for $X < k$, the net change vector is positive, while if $X > k$, the net change vector is negative. The vector field looks like Figure 1.36.

Figure 1.36: The vector field for the logistic equation, $X' = rX(1 - \frac{X}{k})$, with $r = 0.2$ and $k = 100$.

Exercise 1.5.1 If $X' = 0.3X\left(1 - \frac{X}{500}\right)$, what change vector is associated with the point $X = 90$? With $X = 600$?

Exercise 1.5.2 Sketch the vector field for $X' = 0.1X$.

Change Vectors in Two Dimensional Space

In 2D, the state space is a two-dimensional vector space $X \times Y$, which is the space of all pairs (X, Y). The general form for a model in two variables is that X' depends on the full state, that is, on both the X and Y values; Y' also depends on the two values (X, Y). We write

$$X' = f(X, Y)$$
$$Y' = g(X, Y)$$

Let's look at our spring model

$$X' = V$$
$$V' = -X$$

The change vector at the point $(X, V) = (1, 1)$ is $(X', V') = (1, -1)$. So we draw the change vector $(1, -1)$ at the point $(1, 1)$. Similarly, the change vector at the point $(1, -1)$ is the vector $(-1, -1)$, the change vector at the point $(-1, -1)$ is the vector $(-1, 1)$, the change vector at the point $(-1, 1)$ is the vector $(1, 1)$; see Figure 1.37, left.

If we draw many such change vectors, the picture looks like Figure 1.37, right. We can begin to guess what the overall motion is going to be by looking at the change arrows.

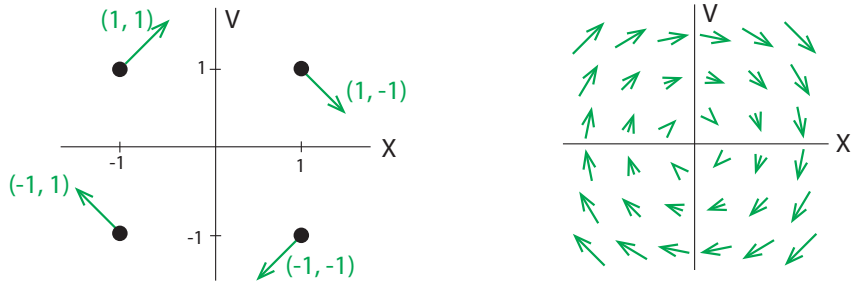

Figure 1.37: Left: Four representative change vectors (green) for the simple mass-spring model, drawn on the (X,V) state space. Right: plotting many change vectors gives us a sense of the dynamics of the system.

Exercise 1.5.3 What would the vector field for the Romeo–Juliet model look like?

Now let's consider a shark–tuna model,

$$T' = 0.5T - 0.01ST$$
$$S' = 0.005ST - 0.2S$$

When $T = 10$ and $S = 10$, we have $T' = 0.5 \times 10 - 0.01 \times 10 \times 10 = 4$ and $S' = 0.005 \times 10 \times 10 - 0.2 \times 10 = -1.5$, so the change vector associated with the point $(10, 10)$ is $(4, -1.5)$. This means that at the point $(10, 10)$, T is increasing at a rate of 4 tuna per unit time, and S is decreasing by 1.5 sharks per unit time.

The change vector is a two-dimensional vector (2D) assigned to each point in the 2D space. So for each state point (X, Y), we can calculate the change vector (X', Y') by computing $\Big(f(X, Y), g(X, Y)\Big)$. This change vector belongs to the 2D tangent space (X', Y') (Figure 1.38).

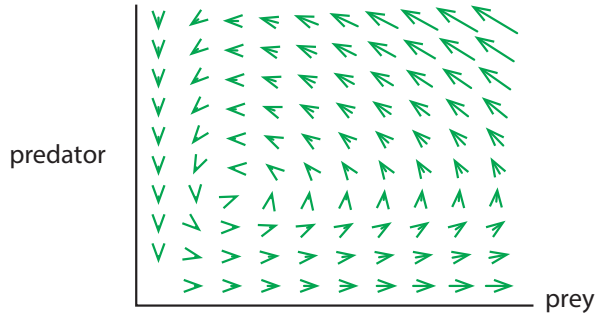

Figure 1.38: The vector field of the Lotka–Volterra predation model $T' = 0.5T - 0.01ST$, $S' = 0.005ST - 0.2S$.

It is easy to plot vector fields on state space using SageMath. Figure 1.39 shows a SageMath output for the vector field

$$X' = 0.9X - 0.5Y$$
$$Y' = 0.1X + 0.8Y$$

```
>>> vector_field(X, Y) =
        (0.9*X - 0.5*Y, 0.1*X + 0.8*Y)
>>> p = plot_vector_field(vector_field,
        (X, -2, 2), (Y, -2, 2),
        frame=False,
        color="green",
        plot_points=10,
        axes_labels=("$X$", "$Y$") )
>>> show(p, aspect_ratio=1, figsize=5)
```

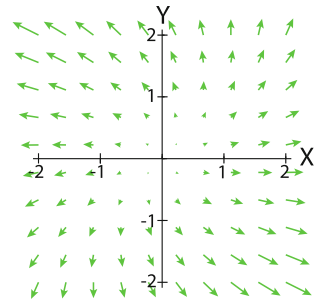

Figure 1.39: SageMath code to produce a vector field and output.

Exercise 1.5.4 If $X' = Y$ and $Y' = X$, what change vector is associated with the point $(3, 5)$?

Exercise 1.5.5 Find the change vector associated with the point $(T = 75, S = 75)$ in the Lotka–Volterra predation model $T' = 0.5T - 0.01ST$ and $S' = 0.005ST - 0.2S$.

Further Exercises 1.5

1. Pick two points on this vector field. For each one, sketch a time series plot describing the system's dynamics and describe them verbally.

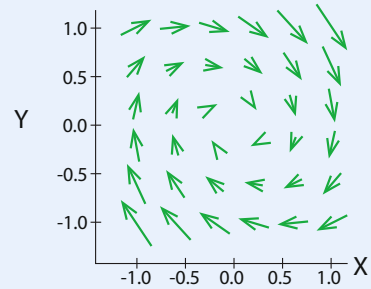

2. Romeo and Juliet are in a relationship. R represents Romeo's love (or if negative, hate) for Juliet, and J represents Juliet's love or hate for Romeo. Suppose that both Romeo's and Juliet's feelings are affected by both their own and the other person's feelings in exactly the same way:

$$R' = aR + bJ$$
$$J' = aR + bJ$$

Let $a = 0.5$ and $b = 1.25$. Plug in numbers to sketch the vector field of this system (include your calculations). Then, describe its behavior.

3. In SageMath, vector fields can be easily plotted with the `plot_vector_field` command. For example, the vector field in Figure 1.39 on the previous page was plotted with the command

```
>> var("x, y")
>> plot_vector_field([0.9*x-0.5*y, 0.1*x+0.8*y], (x, -10, 10), (y, -10, 10)
    , axes_labels=["x", "y"])
```

Redo the shark–tuna vector field (green arrows in Figure 1.38) and the spring with friction vector field

$$X' = V$$
$$V' = -X - V$$

using `plot_vector_field`. Make sure to use a reasonable state space and label the axes correctly.

4. Zebras and wildebeest compete for food on the Serengeti Plain. If Z and W represent the zebra and wildebeest population sizes, the equations representing the population dynamics might be

$$W' = W(1.05 - 0.1W - 0.025Z)$$
$$Z' = Z(1.1 - 0.05Z - 0.2W)$$

Sketch the vector field for this system (include your calculations) and describe what happens to each population as time passes.

5. You can use `plot_vector_field` in a SageMath interactive just as you would use the regular `plot` command. Consider the Romeo–Juliet model in which each person responds only to the other's feelings: $R' = aJ$ and $J' = bR$. (The parameters a and b can be either positive or negative.) Create an interactive that plots the vector field and lets you manipulate a and b. Then, describe how the system behaves at various parameter values. *(Hint: You may find it helpful to organize your observations in a table.)*

1.6 Trajectories

Trajectories in State Space

The picture we have developed so far is that the *state* of a system is a *point* in state space, and changes in the system are represented by *change vectors*, with a change vector assigned to each point in state space. In other words, state space is paved with change arrows at every point. This is the vector field.

We also said that change is movement through state space. Suppose you start at the white circle in Figure 1.40. This is called an *initial condition*. Now imagine that as the state point moves through state space, it leaves a trail behind it. This trail tells us the history of the system—all the points the system has visited. This is the red curve in the figure. This curve is called the "solution curve" or "integral curve" or just "*trajectory*."

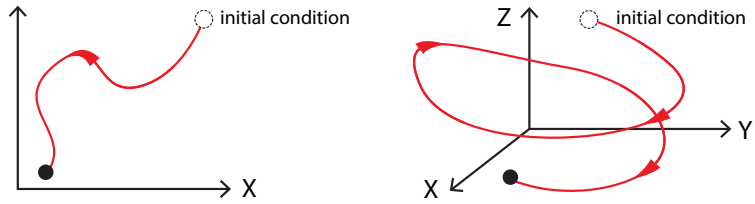

Figure 1.40: Trajectories in two- and three-dimensional state space.

The trajectory arises by following the change arrows of the vector field at every point. Let's use the shark–tuna vector field as an example, and let's start at an initial condition indicated by the black dot at the lower right (Figure 1.41, left). Then the state point, following the change arrows, will move up and to the left, then sharply down, and then to the right.

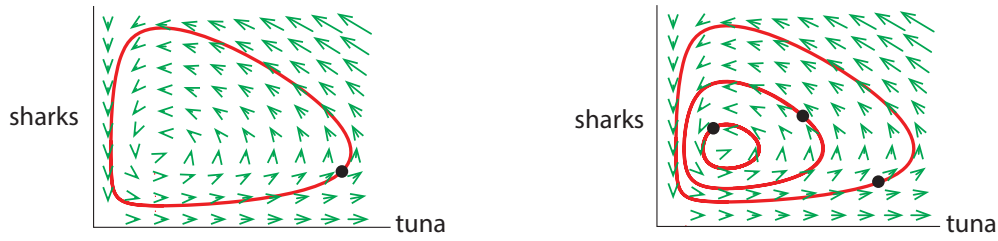

Figure 1.41: Left: One trajectory for the Shark-Tuna model, starting from the initial condition at the black dot. Right: Several different initial conditions (black dots) give rise to distinct trajectories. In this system, the overall behavior depends strongly on the choice of initial condition.

Indeed, if you look at the vector field, you can basically *see* how the state point (black dot) is going to move. It seems to be "following" the green change vectors everywhere. (In the next section, we will define exactly what it means to be "following" the green vectors.) If we choose several different initial conditions for the shark–tuna model, we see that each initial condition gives rise to a distinct trajectory (Figure 1.41, right).

Trajectories are curves through state space that tell us everything about where the system has been (although not how fast it traveled). They are very powerful tools, but learning to interpret trajectories takes some time. We will approach it step by step, with lots of examples.

Drawing Trajectories

Even in one dimension, the concept of trajectory makes sense. Think of the state space of a hot cup of coffee in a cooler room. Let's say we care only about the temperature of the coffee, so the state space is one-dimensional (and nonnegative if we use the absolute, or Kelvin, scale). Then, if the coffee starts off at a hotter temperature than that of the room and subsequently cools off, the state point will have moved from the higher temperature T_1 to the lower temperature T_2 (see Figure 1.42).

Figure 1.42: The trajectory of the temperature of a cooling cup of coffee.

For a more biological example, consider a population that either increases or decreases until it reaches a stable level. Two possible trajectories for such a system are shown in Figure 1.43. Figure 1.44 shows the time series for these two trajectories.

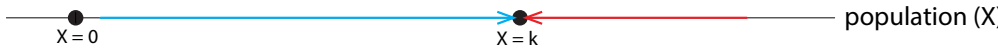

Figure 1.43: Two trajectories of population growth (blue) and decline (red).

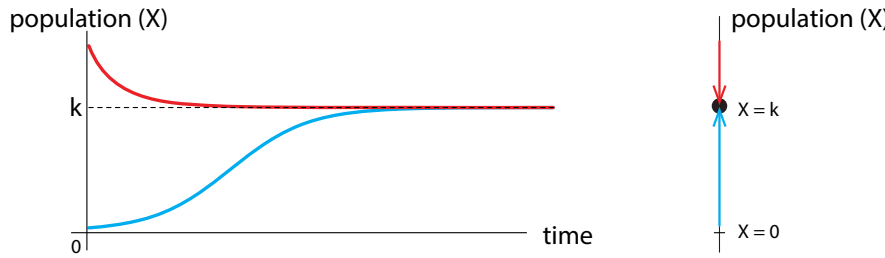

Figure 1.44: Time series corresponding to the two trajectories in Figure 1.43.

Trajectories in 2D Imagine a person who earns a high salary and has correspondingly high expenses. This person then loses most of their income but maintains the same high level of spending as before. Of course, this can only go on for so long, so eventually the person's expenses drop. However, their income then starts to increase, as do their expenses. The upper panel in Figure 1.45 plots the person's income and expenses over time.

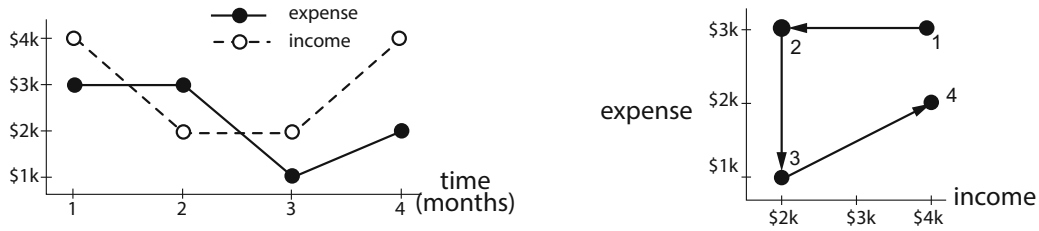

Figure 1.45: Time series and state space trajectories of income–expense dynamics.

We can also graphically depict this story in state space. Let's say that the person's income and expense level together define their state. We can then portray the person's states in income–expenses space, as shown in Figure 1.45. At time 1, which corresponds to the first point on the time series plot, income and expenses are both high. At time 2, income is low, but expenses are still high. At time 3, income and expenses are both low. Finally, at time 4, both income and expenses are intermediate.

For another example, consider a basketball game between the University of X and Y State, as shown in Figure 1.46.

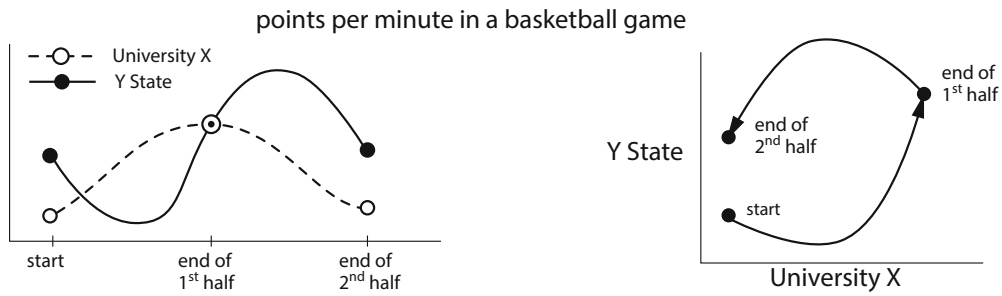

Figure 1.46: Time series and state space of a hypothetical basketball game.

Team Y starts off with a slight advantage but then declines while X's scoring rate increases. Both teams then score more and more points, until at the end of the first half, both have very high scoring rates. X's scoring rate then declines throughout the second half, while Y's increases and then declines but remains fairly high.

From time series to trajectories and back again We will use both time series and state space trajectories heavily. We are accustomed to looking at time series, so we know how to interpret them. But state space trajectories carry critical information about the dynamical system, and it is useful to develop the skill of going back and forth between time series presentations and trajectories in state space.

To go from trajectories to time series, imagine that you are tracing the trajectory with your finger. If you want the X time series, ask yourself, "is X getting larger or smaller?" as your finger traces the trajectory. Then sketch that as a time series.

For example, in the frictionless spring (Figure 1.47), we start at the black dot, which is a negative value of X. As the state point traces clockwise, X declines into more negative values until it reaches its lowest value at 9 o'clock. Thereafter, X steadily and smoothly increases until its maximum at 3 o'clock, whereupon it starts to decline, completing the cycle.

In the spring with friction (Figure 1.48), the state point traces out a spiral: a circular motion with ever-decreasing amplitude. This produces a time series called a *damped oscillation*.

Alternatively, imagine a strip chart recorder (like the one in a seismometer) with a long strip of recording paper passing continuously through the X-V plane and recording the X value of the state point at each moment in time (Figure 1.47, Figure 1.48, and Figure 1.49).

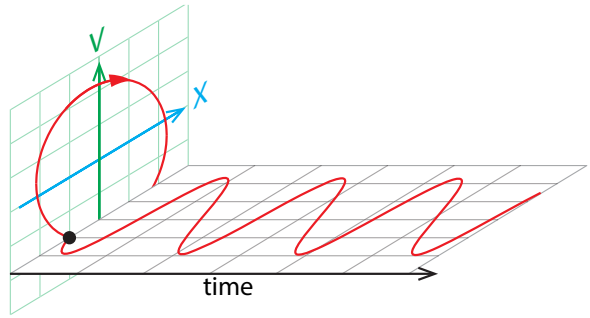

Figure 1.47: A circular trajectory, such as generated by the frictionless spring, produces a smoothly changing periodic time series.

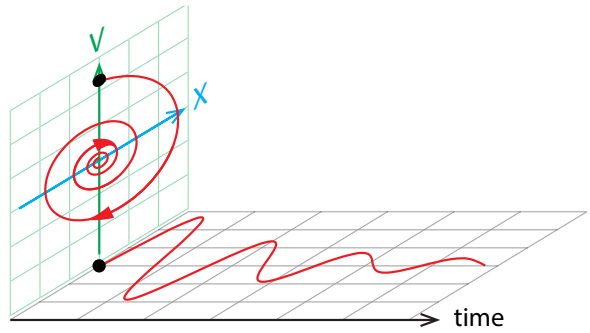

Figure 1.48: If we consider the spring with friction, the resulting trajectory spirals in. This produces a periodic function of time with constantly decreasing amplitude.

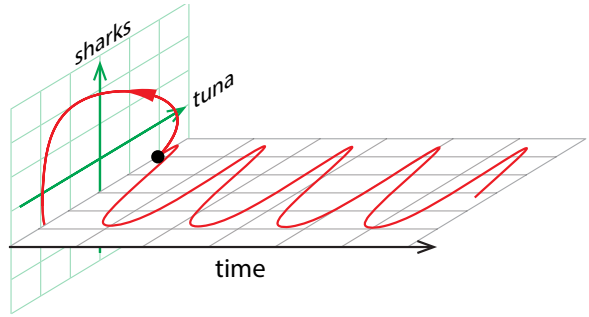

Figure 1.49: Nonround trajectories, as seen, for example, in the shark–tuna model, produce periodic time series with different waveforms.

In general, it is very useful to view system behavior as a trajectory through a state space. As you will see throughout this text, this approach allows us to classify types of behavior and relate them to each other in ways that would be much harder if all we had were time series graphs.

Exercise 1.6.1 Draw a time series graph that corresponds to this trajectory of temperature and day length over the course of a year.

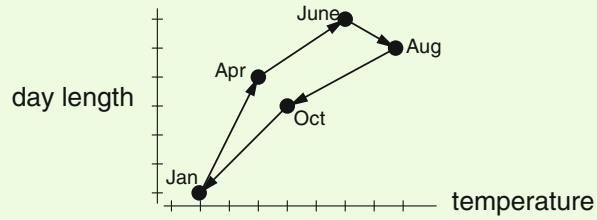

Exercise 1.6.2 Draw a trajectory corresponding to this time series of a person's height and weight over the course of their life.

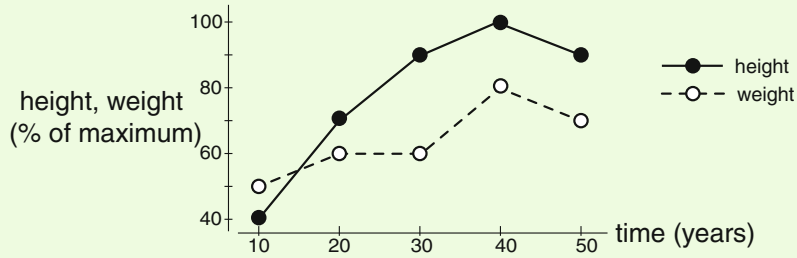

Glucose and insulin Consider the dynamics of glucose (G) in the body, as it is metabolized with the help of the hormone insulin (I). We can make (I, G) space, the space of all pairs (I, G), where G is the blood glucose level and I is the blood insulin level. What happens after a meal (the black dot)? First glucose goes up quickly, then insulin starts to rise, which causes glucose to decline. We can represent this as a trajectory through (I, G) space.

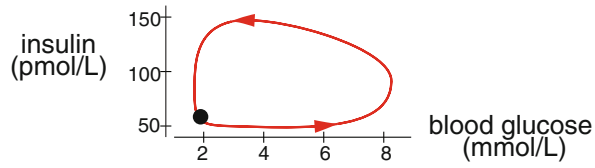

Chemistry In chemistry, we are usually interested in the concentrations of chemicals and how they change over time. For example, we could describe a simple chemical reaction with two state variables, the sodium ion concentration $[Na^+]$ and the chloride ion concentration $[Cl^-]$. Then, as they combine to make NaCl, their concentrations would change:

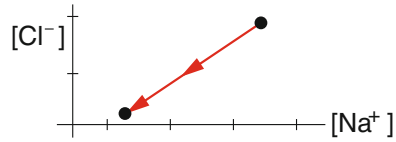

Romeo and Juliet Suppose that Romeo and Juliet go through some changes in their relationship, and go from (Love, Love) to (Love, Hate) to (Hate, Hate). If we plot that as a trajectory through R-J space, it looks like this:

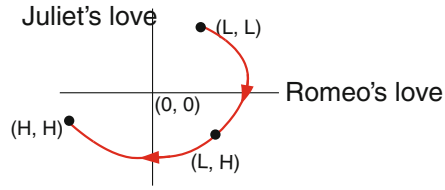

Mechanics of springs The trajectories of frictionless springs look like concentric circles, with each orbit corresponding to a different initial condition:

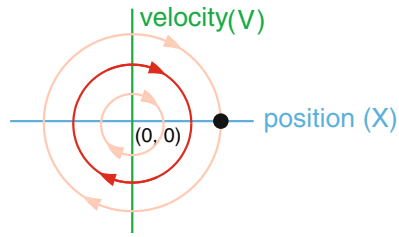

3-Dimensional Systems

Wolves, sheep, and grass As another example, in Chapter 5, we will study a system with three species in which Z eats Y, and Y eats X. There are many examples of such chains, including wolves, sheep, and grass.

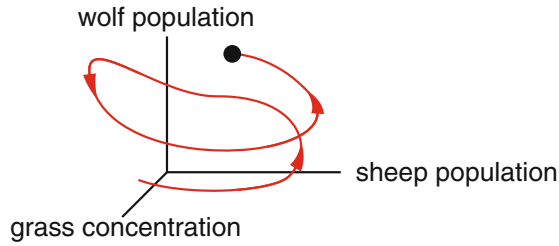

Epidemiology Epidemiology is the study of how diseases affect populations. Typical state variables in epidemiology are S, the number of susceptible individuals, and I, the number of infected individuals. A third variable in epidemiology models is often R, standing for recovered individuals, or, in the Anderson–May model of HIV transmission (see page 40), it could be A, the number of symptomatic AIDS patients.

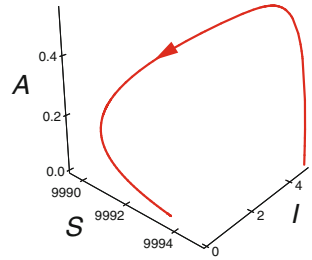

Insulin, glucose, and glucagon High levels of glucose in the bloodstream (for example, after a meal) cause the pancreas to release insulin. But low levels of glucose cause the pancreas to release another hormone called glucagon. So the state of this system is represented as a point in (I, G, A) space, where I is the concentration of insulin, G is the concentration of glucose, and A is the concentration of glucagon.

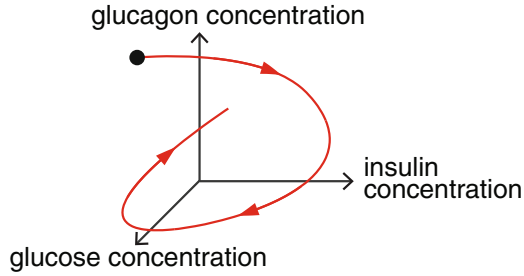

Systems with Four or More Dimensions

The neuron In their Nobel Prize–winning work, Alan Hodgkin and Andrew Huxley showed that the firing of a neuron can be represented by four variables standing for voltage, a current called I_{Na} carried by sodium ions, a current carried by potassium ions (I_K), and a current called "I-leak" that we now know is carried primarily by potassium ions (I_L). Therefore, the state space for the neuron is (V, I_{Na}, I_K, I_L) space, and the course of the neuron's firing is represented as a curve through 4-dimensional (V, I_{Na}, I_K, I_L) space.

Food webs We have already discussed ecological models with two or three predator and prey species. However, real ecosystems have many more species than that. Their complex feeding interactions create what's called a *food web*. In models of food webs, the state variables are the population sizes of different species, and there can be tens to hundreds of them.

The **state** of a system at a given time is a *point* in $\mathbb{R}^n$, $\{(X_1, X_2, \ldots, X_n)\}$. **Change** in a system over time is a *curve* or *trajectory* through state space $f : \mathbb{R}_+ \to \mathbb{R}^n$.

The State Space Trajectory View

Looking at trajectories in state space gives us insights into system dynamics that can't be gotten by looking at separate time series plots.

We said earlier that the results of intervening in a feedback system can be counterintuitive (see Figure 1.9 on page 7), and that the system's response to an intervention can depend on the intervention's magnitude and timing.

Then how can we predict what kind of intervention will produce what kind of results? The answer is that the state space trajectory gives this insight in a "master view." We will illustrate this with the Lotka–Volterra predator–prey model (again with the caveat that there are better models; see the Holling–Tanner model in Chapter 4).

Suppose the system is at state point #1 (Figure 1.50). We are considering only "predator removal" interventions, which means that we are going to move the state point vertically downward by a given amount. It is obvious from the trajectory view that if we start at point #1 and instantaneously remove a small number of predators that takes the system to state point *A*, then the system will go to the new trajectory containing point *A* and will orbit in a smaller trajectory, thus decreasing both shark and tuna populations.

But if we remove a large number of predators, thereby taking the system to state point *B*, then the resulting trajectory is a larger orbit, and the predator population will rebound to a higher level than before the intervention.

And if the system is at state point #2, then *any* predator removal will result in a rebound to a higher peak predator population.

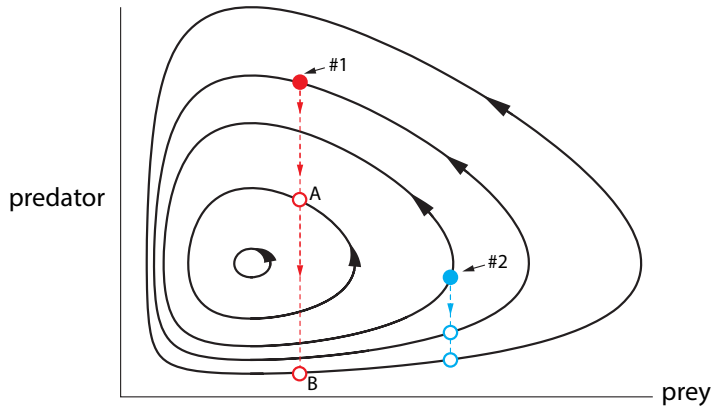

Figure 1.50: Response of the predator-prey (Shark-Tuna) system to perturbations depends on the strength and timing of the perturbation. The outcome of a perturbation is to place the state point on a new trajectory, whose amplitudes may be higher or lower than before.

Thus, the response of feedback systems to intervention, which can be difficult to understand by looking only at the time series, can be easily grasped by looking at state space trajectories.

Vector Fields, Trajectories, and Determinism

In the discussion above, we said that in a vector field, a change vector is associated with every point in state space. Knowing the point, we calculate the unique change vector associated with

it. Since this relationship is completely unambiguous, it is a function. Since the vector field links each point in the state space to a change vector in the tangent space, we can write

$$\text{vector field} : \text{state space} \to \text{tangent space}$$

Furthermore, if our system has n variables, a change vector must have n components—one for each variable. Therefore, change vectors live in $\mathbb{R}^n$, and we can write

$$\text{vector field} : \mathbb{R}^n \to \mathbb{R}^n$$

A **vector field** is a function $V : \mathbb{R}^n \to \mathbb{R}^n$ that assigns change vectors to state vectors.

The fact that vector fields are functions turns out to have important implications for dynamics. Since a vector field is a function, there is exactly one change vector associated with each point in state space. This makes it impossible for trajectories to cross! Why? Because trajectories always follow change vectors.

Suppose two trajectories crossed and then went off in different directions. Then, as illustrated in Figure 1.51, there would have to be two change vectors at the point of intersection—one that the first trajectory followed and one that the second trajectory followed. But then the vector field wouldn't be a function. Therefore, trajectories can't cross. As we will see, this is a powerful constraint that tells us a lot about dynamical behavior.

Exercise 1.6.3 In this situation, why would the vector field not be a function?

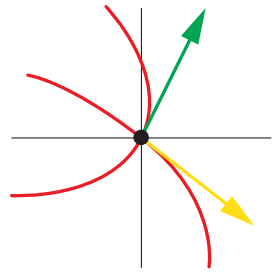

Figure 1.51: What a vector field would look like if two trajectories crossed.

There's more. We saw that trajectories cannot cross. As we will see in Section 1.7, they also cannot touch. The uniqueness of trajectories at each point means that if we know a system's state at any time, we can find its trajectory for all time.

Further Exercises 1.6

1. This trajectory shows the hours a student studied per week during a quarter and that student's GPA for that quarter. Describe what happened and sketch the appropriate time series plots.

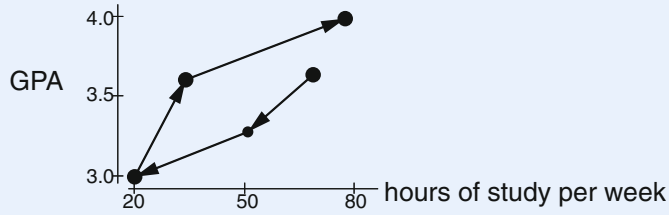

2. Sketch the time series of a two-variable system whose trajectory is a single point.

3. This trajectory was generated by a simple model of an oscillator. Sketch the time series matching it.

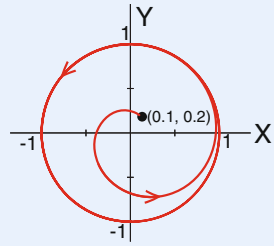

4. This time series graph describes the spread of an infectious disease. Sketch the trajectory corresponding to this time series.

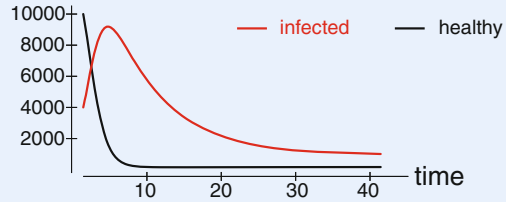

5. Consider the time series graphs for the lynx and snowshoe hare populations in Figure 1.1 on page 1, which we repeat below.

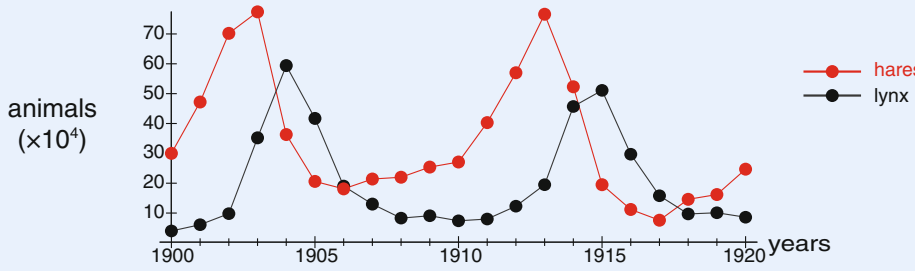

Sketch an approximate trajectory for this system. *(Hint: Pay attention to the points at which populations go from increasing to decreasing and vice versa.)*

6. Come up with your own example of a two-variable system that changes over time. Describe it verbally and draw a time series graph and matching trajectory.

7. How are a trajectory and a time series graph different? In particular, what are the axes of each? Which one can reasonably be drawn on top of a vector field?

1.7 Change and Behavior

In the previous section, we drew a number of trajectories. These are the red curves that trace out a system's behavior. We saw red trajectories in the shark–tuna model (Figure 1.41) and in the spring (Figure 1.47, Figure 1.48).

How did we get these red curves? We claimed that the red curve is "following" the change arrows, but you might ask two questions:

Q1: Does that red curve really exist? Is there really a single trajectory through a given point that everywhere follows the change arrows?

Q2: Can we figure out the *equation* for the red curve from the *equation* for the vector field?

The answers to those questions are

A1: Yes, almost always.

A2: No, almost never.

There is a theorem that answers question Q1. It says that if our differential equations are well behaved (and everything we will see in this course is well behaved; the basic idea is that the functions can't change too fast), then *there is a unique curve through any given point that "follows" the change vector at that point.*[8] So that red curve truly exists; it is "out there."

How do we find that red curve? We have the model as a vector field

$$V : \{\text{states}\} \to \{\text{changes in state}\}$$

We would love to go from a state to its change of state and then to a new "next" state. Now, a change vector tells us how the system would change if it followed the change vector one whole time unit (one year, one day, one second, etc.). It therefore seems natural to take a state, add the change vector associated with that state, and use the result as our next state.

But we have a problem: there is no real "next" state. Recall that when we talk about the value of a state, X, we really mean its value at some particular time t. If our initial time is $t = 0$ and we take $t = 1$ to be the "next" time, someone could point out that $t = 0.5$ is between 0 and 1, so $t = 1$ can't be the next time. Similarly, 0.25 is between 0 and 0.5, so 0.5 can't be the next time either. In fact, there are infinitely many numbers between any two real numbers, so we would have to update the state infinitely often, which is clearly impossible.

This was the problem that faced Isaac Newton in the late 1600s. He suspected that the force of gravity acting between the Sun and the Earth was causing the movement of the Earth. He also knew that forces change velocities, that is, produce accelerations, and so he could say

$$\begin{aligned}&\text{old position} \to \text{force} \to \text{acceleration} \to \text{change in velocity}\\ &\to \text{change in position} \to \text{new position} \to \text{new force} \to \ldots\\ &\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{(and so on)}\end{aligned}$$

[8] This is called the Picard–Lindelöf theorem, or the fundamental theorem on the existence and uniqueness of solutions to ordinary differential equations.

But how to make that update *at every time point*? This is where calculus, which you will learn about in the next chapter, comes from. Say you are at the state X_0 and so follow the change vector out of X_0, which is X_0'. *But for how long do you follow the change vector?* If you follow it for a second you are wrong. If you follow it for a tenth of a second, you are still wrong. This is because long before that tenth of a second was up, you were already at an X point different from the one you started with, and that point has its own change vector. So even a brief moment into that tenth of a second, you were already using the wrong change arrow.

Let's call the amount of time you follow the change vector Δt. Newton made Δt smaller and smaller, and he saw that in the case of the gravity vector field, if he let Δt approach zero, he could actually calculate the equation for the red curve. That's called calculus, and we will discuss this in more detail a little later.

Calculus, meaning letting Δt go to 0, using the concept of infinitesimal limits and then figuring out the equation for the red curve, is great, when it can be done. *But it can almost never be done!* It can't be done for the shark–tuna model, for example. Remember those curves of $S(t)$ and $T(t)$ (Figure 1.2 on page 2)? **The equations for those curves are unknown.** In fact, virtually none of the models we encounter in biology have a solution curve whose equation can be found.

So how are we able to plot these graphs? The answer is called *Euler's method*.

Taking Small Steps

Euler's method consists in making Δt very small but not zero. Starting at our initial point, we will follow its change vector for a very short time, specifically Δt. We then find the change vector associated with the new point and follow it for the same very small time interval. *Doing this over and over gives us a good approximation to the red curve*, especially if we choose our Δt to be very small.

Euler's Method in One Dimensional Space

Suppose we are dealing with a one-dimensional state space X and a differential equation $X' = f(X)$. Let's start from an initial condition X_0. Then the change vector at X_0 is $f(X_0)$. Since we're following this change vector for only Δt time units, the actual change is not $f(X_0)$ but $\Delta t \cdot f(X_0)$. (Recall that we can multiply vectors by constants.) To get the new state, we just add this amount to the old state:

$$\text{new } X = \text{old } X + \Delta t \cdot X'$$

Applying this procedure over and over is called Euler's method.

For example, suppose X is the size of an animal population and the growth rate of the population is modeled by $X' = 0.2X$. Suppose we start with a hundred animals, so $X_0 = 100$. Let's choose a nice small step size, say $\Delta t = 0.01$. Then

$$\begin{aligned}\text{new } X &= X_0 + 0.01 \cdot f(X_0)\\ &= 100 + 0.01 \cdot 20\\ &= 100.2\end{aligned}$$

We will now call the new X above X_1, and one step of Euler's method is complete. To start the second step, X_1 becomes the old$X$, and we have

$$\begin{aligned}\text{new new } X &= X_1 + 0.01 \cdot f(X_1)\\ &= 100.2 + 0.01 \cdot 20.04\\ &= 100.4004\\ &= X_2\end{aligned}$$

Exercise 1.7.1 Compute X_3.

Exercise 1.7.2 Use Euler's method to compute two approximate trajectories for the logistic growth vector field $X' = 0.05X(1 - \frac{X}{100})$.

Euler's Method in Two Dimensional Space

The geometric meaning of Euler's method becomes especially clear when we look at the 2D case. Now our state variables are X and Y, and our differential equations have the form $X' = f(X, Y)$ and $Y' = g(X, Y)$. These equations create a vector field on $\mathbb{R}^2$.

Now we write Euler's method in two parts:

$$\begin{aligned}\text{new } X &= \text{old } X + \Delta t \cdot X'(\text{old } X,\ \text{old } Y)\\ \text{new } Y &= \text{old } Y + \Delta t \cdot Y'(\text{old } X,\ \text{old } Y)\end{aligned}$$

Let's simulate the shark–tuna model with Euler's method. If we set all the parameters to 1, the equations are

$$\begin{aligned}S' &= ST - S\\ T' &= -ST + T\end{aligned}$$

Let's take as our initial condition (S_0, T_0) the point $(2, 3)$ in S-T space, and let's take $\Delta t = 0.1$. We first calculate the change vector at this state: $S' = 2 \cdot 3 - 2 = 4$ and $T' = -2 \cdot 3 + 3 = -3$. So the change vector is $(S', T') = (4, -3)$. Then the first iteration of Euler's method is

$$\begin{aligned}\text{new } S &= 2 + 0.1 \cdot 4 = 2.4\\ \text{new } T &= 3 + 0.1 \cdot (-3) = 2.7\end{aligned}$$

Exercise 1.7.3 Compute the next values of S and T.

When we do that for many iterations and we keep our Δt small, the resulting vectors, tip to tail, approximate the true red curve very well.

Indeed, we have been showing you a number of red trajectory curves. Where did we get those curves? In the case of the shark–tuna vector field, for example, the equation for the red curve is unknown. So how could we draw it? The answer is that we aren't really drawing the red curve; what we are doing is drawing a blue broken-line approximation with a Δt that is so small that the jagged approximation looks smooth to the eye (Figure 1.52).

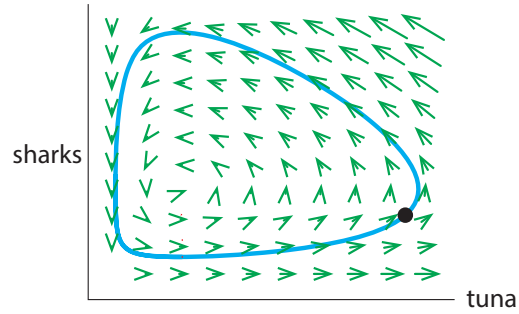

Figure 1.52: Euler's method approximation to the shark–tuna model. The short blue straight lines of Euler's method are too small to be seen here. The resulting trajectory of straight lines looks like a smooth curve.

The geometric picture in Figure 1.53 is the clearest way of seeing what is going on. We are approximating the smooth red curve by the jagged blue line. (There is a mathematical theorem called the shadowing lemma, which says that as Δt gets smaller and smaller, the blue jagged line gets closer and closer to a true red curve, possibly from a slightly perturbed initial condition.)

Euler's Method

1. Start from the point X_0.
2. Evaluate X' at X_0. We will call this X_0'.
3. Multiply the change vector X_0' by the small number Δt.
4. Add $\Delta t \cdot X_0'$ to X_0, and call the result X_1.
5. Repeat steps 1 through 4 for the point X_1 to get X_2. Then repeat for X_2 to get X_3, etc.

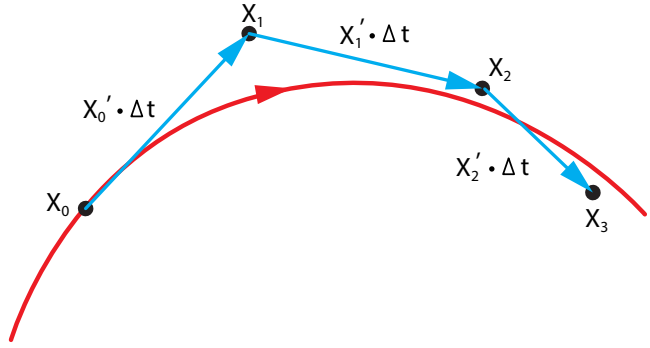

Figure 1.53: Euler's Method. The red curve is the true trajectory of the system. Beginning at the point X_0, one step of Euler's method, with a step size Δt, (blue arrow), takes the system to the point X_1. A second step of Euler's method, from the point X_1, takes the system to the point X_2, and a third step takes the system to the point X_3, forming an approximation to the red curve.

Generating a time series or trajectory from a model and an initial condition is called *simulating* or *numerically integrating* the model. There are other simulation methods that approximate trajectories more accurately than Euler's method. The math behind these methods is slightly more complicated and need not concern us. However, some of these methods are built into SageMath, and we will use them extensively later on.

You may have noticed that the repetitive nature of Euler's method makes it ideal for computers. Indeed, large-scale numerical integration is unpleasant without a computer, which is why dynamical simulation is heavily computer-dependent. However, it is perfectly possible to do numerical integration by hand, and there are many famous examples of this.

Numerical integration without computers

In the days before electronic computers, numerical integration was done by hand. For example, the return of Halley's Comet in 1758 was predicted by numerical integration of Newton's equations by hand. Integration by hand was also used to calculate artillery trajectories in World War I. See the excellent book *When Computers Were Human*, by David Alan Grier (Grier 2013).

Even in the late 1940s and 1950s, numerical integration was still frequently done by hand. Hodgkin and Huxley used it to do their simulation of the firing of a neuron, for which they received the Nobel Prize. In the early years of the US space program, human computers, many of whom were African-American women with math degrees but limited employment options, worked in aeronautical engineering at the National Advisory Committee on Aeronautics, which later became NASA. The book *Hidden Figures* by Margot Lee Shetterly and the movie based on this book tell their story (Shetterly 2016).

Further Exercises 1.7

1. The rate of change of a mouse population is given by the differential equation

$$N' = 0.5N\left(1 - \frac{N}{1000}\right)$$

   The population at $t = 0$ is 400. Using Euler's method with a step size of 0.1, find the (approximate) population at $t = 0.3$.

2. The growth rate of a hunted lion population, L, is given by the differential equation

$$L' = 0.1L\left(1 - \frac{L}{100}\right) - 0.2L$$

   The current population is 80 lions. Using Euler's method with a step size of 0.1 years, find the (approximate) population 0.2 years later.

3. A disease is spreading in a population. We will model the number of susceptible individuals (S) and infected individuals (I) using the differential equations

$$S' = 0.2I - 0.05SI$$
$$I' = -0.2I + 0.05SI$$

   Suppose we start with 98 susceptible individuals and 2 infected ones. Use Euler's method with a step size of 0.1 weeks to determine the approximate numbers of susceptible and infected individuals at $t = 0.2$ weeks.

4. Briefly describe the advantages and disadvantages of using a very small step size in Euler's method.