

On two occasions I have been asked, “Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?” I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question.

Charles Babbage, English inventor of the first mechanical computer

Chapter 4

Fundamentals of functions

We have seen a number of relationships so far: the relationship between distance and time of a moving object, for example, or the relationship between a number’s position in a sequence and that number’s value. A *function* is a certain kind of relationship that is of key importance to us in algebra and, indeed, throughout mathematics. In this chapter we’ll get an overview of functions and their different representations. Then, in the rest of the course, we will look closely at three specific kinds of functions.

4.1 Mathematical relationships

Number machines

// startup exploration

A number machine accepts numbers as input and produces numbers as output. When machine A receives the number 8 as input, it produces the output 28. When this machine receives -12 as input, the output is -42 .

Machine B has a different mechanism for producing output values. When machine B receives the number 8 as input, it outputs 15. When it receives -12 as input, the output is 23.

What do you suppose each machine will output when given the number -7 as input? Write a sentence explaining how you believe each machine works.

The number machines described in the startup exploration relate certain input values to certain output values. Mathematically speaking, a number machine defines a *relation*.

Relation

A **relation** defines how certain numbers (or other objects) are connected to other numbers (or objects). We can think of a relation as a collection of ordered pairs (x, y) , meaning “ x is related to y ”.

We might think of a number machine as a set of pairs of numbers: (input, output). Machine A in the startup exploration is defined by the pairs $(8, 28)$ and $(-12, -42)$. We sometimes say that “machine A maps 8 to 28” and that “-12 is mapped to -42”.

One possible explanation is that the machine multiplies the input value by 3.5 to produce the output value. In this case, we’d expect the input -7 to produce the output -24.5 . In other words, $(-7, -24.5)$ is a third point associated with machine A.

For machine B we are given the points $(8, 15)$ and $(-12, 23)$. One, somewhat complicated, explanation is that the machine doubles the input, takes the absolute value of the result, and then subtracts 1. In this case, we’d have $(-7, 13)$ as another input-output pair for machine B.¹

We can define relations that compare things other than numbers. For example, we might define the “name-the-mother” relation that accepts a person as input and which gives that person’s biological mother as output. Or, we might define the “name-the-pet” relation, which accepts a human as input and outputs other animals.

Note that there is a difference between these two relations we have just defined. Not all people own pets, so some inputs to the name-the-pet machine have no output. Plus, some people own multiple pets, so some inputs to the name-the-pet machine have multiple outputs. On the other hand, the name-the-mother machine is guaranteed to always produce a unique output: every person has a biological mother, and it is impossible for a person to have more than one biological mother!²

We saw something similar in section 3.5, where certain graphs could be drawn in a way that would suggest an impossible situation. For example, we drew a graph (shown in fig. 4.1) that suggested Yearleigh’s submersible could at three different depths simultaneously. This is as ridiculous as the idea that a person could have more than one biological mother!

Relationships which follow this very natural rule – any given input produces a unique output – have a special status in algebra. They’re called **functions**.

¹ There are other rules that we could write based on the two data points we were given, and perhaps you thought of different rules than the ones we describe here. Can you find alternative explanations for how each machine works? Be creative!

² We don’t deny that a family with two mothers is still a family! It is true that a person can have multiple mothers in a legal or emotional sense. . . but a person always has exactly one *biological* mother.




Figure 4.1: A graph of depth over time? Impossible!

Function

A **function** is a special type of relation in which the ordered pairs have the following property: each x -value is paired with one, and only one, y -value.

The “one and only one” requirement is what makes a function special. The graph in fig. 4.1 violates this requirement: certain x -values have more than one associated y -value (look at x values between 15 and 45 minutes). So, this graph does not depict a function.

The sequences we studied in chapter 2 are functions: it would be silly to suggest that a sequence could have more than one value as its third term. Depending on how we think about number machines, we might expect that a machine would act predictably and always produce the same output for a given input. A number machine that behaves in this way defines a function.

Representing relations

There are several ways to represent a mathematical relation. We might represent a relation using:

- A graph. This representation gives us a visual of the relationship between input values and output values.
- A table of values. This representation gives us an organized list of input values and their corresponding output values. (Related forms include a *mapping diagram* or just a collection of ordered pairs.)
- An equation. This representation gives us a rule for turning input values into output values, the way a number machine might do.


Since functions hold a special position in the universe of all relations, our first task is to determine whether a given mathematical relationship is, in fact, a function. To do this we have to check it against the definition of what it means to be a function. In other words, we must make sure each x -value corresponds to one and only one y -value.

4.1.1 Graphs of functions

In order for a graph to be a function it must avoid the situation in which multiple y -values stack up above a particular x -value. A handy way of remembering this requirement is that the graph of a function passes the **vertical line test**. The vertical line test states that for the graph of a function, any vertical line drawn on the graph will intersect the graph *exactly once*.

Example 4.1

Which of the graphs below, if any, represents a function? How do you know?




Solution: Only graph A represents a function. We can draw vertical lines on graphs B and C which intersect the graph at more than one point. This means that those x -values have more than one y -value, violating the definition of function.

Note that in graph C, certain sections of the graph pass the vertical line test. There is no partial credit, though. A graph fails the vertical line test if any vertical line (even just one) crosses the graph in more than one point. This happens also with the graph in fig. 4.1. That graph fails the vertical line test for some vertical lines. Therefore the relationship depicted is not a function.

Example 4.2

Which of the scatter plots below, if any, represents a function? How do you know? You may have to look closely!



Solution: Plots A and C both pass the vertical line test, and therefore represent functions. Plot B fails the vertical line test because it fails for the vertical line at $x = 3$.

Note that we don't care if different x -values map to the same y . For example, in plot C, all of the x -values map to the same y -value. That's OK. In other words, a function *can* fail the "horizontal line test" without penalty.³

For example, the name-the-mother function does not pass the horizontal line test because two different inputs (people) can have the same output (mother). In fact, it is not at all uncommon for two different people to have the same mother: they're called siblings!

4.1.2 Functions from points

A scatter plot gives a visual representation of a collection of ordered pairs. If the ordered pairs are presented in an alternative format – say, a table of values – we can use the same reasoning to determine whether the given relation is a function.

³ The horizontal line test *is* a thing, and passing or failing the horizontal line test does have a mathematical meaning. A discussion of this will have to wait until later, however.

Example 4.3

Which of the data sets below, if any, represents a function? How do you know?


Set A		Set B		Set C
x	y	x	y	
-3	4	2	1	$\left\{ \begin{array}{l} (0, 4); \quad (-1, 6); \\ (2, 3); \quad (0, 4); \\ (1, 5); \quad (-2, 5) \end{array} \right\}$
-2	6	1	2	
-1	4	-1	7	
0	2	4	3	
1	0	3	4	
2	-2	1	0	
3	0	0	6	

Solution: Set A is a function. There are no repeated x-values! Set B is not a function. The x-value 1 appears twice and with two different related y-values (0 and 2). This violates the definition of function. Set C is a function. The x value 0 appears twice, but it is mapped to the same value in both cases.

A **mapping diagram** is a similar way of representing a relationship between specific input and output values. Here we draw two regions (represented by a box or an oval). In one region, we list the input values, in the other region we list the output values. Then we draw arrows to show which input maps to which output.

Example 4.4

Which of the mapping diagrams below, if any, represents a function? How do you know?



Solution: The first mapping is a function. Each x-value maps to exactly one y-value. The second mapping is not a function: the value 2 maps to both 7 and 9.

Note again that the problem in diagram 2 is *not* that multiple arrows point to each of the output numbers. It's OK for multiple arrows to go *toward* a certain output value. The problem is that the input value 2 has multiple arrows *leaving* and going to different targets (both 7 and 9).

4.1.3 Equations of functions

The “number machine” metaphor is often helpful when thinking about a function. A number (an x -value) goes into the machine, the machine’s rule works on the number, and another number comes out of the machine (the y -value). Most of the time, when we have a particular rule and a particular input, we get a single, predictable output.

For example, the equation $y = 3x$ defines a function. We can substitute in 5 for x . The function machine multiplies 5 by 3 and outputs 15. This generates the ordered pair $(5, 15)$ and we say that “5 maps to 15”. If we choose other values of x , we can generate more ordered pairs.

Not every equation defines a function, though. In order to be a function, every input value must produce one and only one output value. In Algebra 2 we will study many interesting and important mathematical relationships that are not functions. We will learn, for instance, how to write the equation for a circle – and a circle is not a function. (Can you explain why not?)

Our focus in Algebra 1, though, will be on three main “families” of functions, and our rules will all be of the form “ y equals some expression in terms of x ”.

4.1.4 Function notation

We can thank Swiss mathematician Leonhard Euler for the concept of a function.⁴ He was the first to coin the term “function” and was the first to use a handy way of writing a function, called *function notation*.

Up until now, we have used “ $y = \text{something in terms of } x$ ” to define all of our rules. Euler’s function notation looks a bit different, but it is quite helpful in certain situations. The generic form of function notation is “ $f(x) = \text{something in terms of } x$ ”. In this case, we simply replace y with the notation $f(x)$, which is read aloud as “ f of x ” or “ f as a function of x ”.

In its generic form, the letter f is used to name the function. We use the letter f , of course, because it stands for “function”. And as usual, we use x to represent the independent variable.

What’s nice about this notation is that it allows us great flexibility to use other letters or symbols that can be more descriptive. For example, if we wanted to describe how the height of a bouncing ball changes over time, we might want to let the variable t represent time and then name the function h for height.

So, we could write a function “ $h(t) = \text{some expression in terms of } t$ ”. That’s read “ h of t ”, which does a pretty good job of capturing the idea that we’re describing *height* in terms of *time*.

We might write a function like $d(t) = 60t$ which described distance traveled d as a function of time t . This is exactly the same equation as $y = 60x$, but we’ve changed the names of things to better match the context.

⁴ The surname Euler is German, and so it is pronounced *OIL · er* (and *not* *YOO · ler*).

Evaluation using function notation

Function notation has another convenience. We can use the notation to indicate a specific choice of value for the independent variable. Earlier, we said: “Evaluate the function $y = x^2 - 4$ when $x = 3$.” This was fine, but function notation gives us a way to shorten these instructions.

For example, if we have $f(x) = x^2 - 4$ then we might want to evaluate $f(3)$, which is said aloud “ f of three”. The x in the function notation was replaced by 3, which is exactly what it means to evaluate the function at that value! The notation is telling us to input 3 for x and find the output value. So if $f(x) = x^2 - 4$, then

$$f(3) = (3)^2 - 4 = 9 - 4 = 5.$$

We will use function notation and $y =$ notation through this course. The flexibility of function notation means that it is the preferred way of writing functions in higher level mathematics courses.

4.2 Domain and range

Machines (meaning real, mechanical devices) are physical things that have physical limitations. Number machines (meaning functions) often have similar operational restrictions.

The ins and outs of functions

// startup exploration

Consider the function machines below.

$$f(x) = |x| \qquad g(x) = \frac{1}{x}$$

What sorts of values are allowed as input to each of these machines? What sorts of values will each one produce as output?

A blender is a machine that has limited operating parameters. Anything we put into the blender comes out “blended” in some way. But, of course, there are some things we can’t put in a blender: a piano, for example.⁵ Similarly, we wouldn’t put bread into a blender, hit the button, and expect toast to pop out. That kind of output is associated with a different machine.

4.2.1 Regarding input values

An important component of the definition of any function is a specification about what sorts of numbers it can accept as input. In the language of mathematics, we call the set of all “legal” input values the *domain* of the function.

Domain

The set of all allowed input values to a function. This is the set of all allowed x -values or all possible values of the independent variable.

We have two functions in the startup exploration. The first function, $f(x) = |x|$, is very welcoming: any real number can be used as input to this number machine: whole numbers, fractions, positive numbers, negative numbers, it doesn’t matter. The other function is a little more finicky.

⁵ We’ll pause here to note that “blended” is defined rather loosely: smoothies come out “blended”, ice comes out “crushed”, chickpeas come out “puréed”, magazines come out “shredded”... We consider all of these to be a kind of “blended”. Speaking of magazines, we’ll also mention that there are things which one *can* put into a blender, but *shouldn’t*. Magazines are one example. Another example: your hand.

In chapter 1, we discussed the illegality of division by zero. So, if we try to put 0 into the function $g(x) = \frac{1}{x}$, then we're going to have a problem: $g(0)$ – that's the function g evaluated at $x = 0$ – is undefined. So, our user manual for the function g has to include a note that only numbers other than 0 are allowed as input. We say that the function $g(x)$ has a *limited* or *restricted domain*.

There are other circumstances where we will want or need to limit the domain. For example, a domain can be limited by the context of a problem we are solving. If we write an equation that shows distance as a function of time, it usually makes sense to limit the domain to only include positive time-values. If we have a function which maps “number of pizza toppings” to “price of the pizza”, then there is (probably) an upper limit on the domain: there are only so many pizza toppings available.⁶

A sequence is a function whose domain is limited to the set of natural numbers (the positive integers). We saw this in chapters 2 and 3: a sequence can have a first term and a second term, but no one-and-a-halfth term. We exclude from the domain of a sequence all numbers except $1, 2, 3, 4, \dots$.

We can also define a specific domain for convenience. On an assignment, for example, we might be asked to evaluate the function $y = 4x + 5$ for the x -values $-1, 0$, and 1 . We did something like this in example 3.1. This is another case of limiting the domain.

4.2.2 Regarding output values

In addition to knowing about input values, it is often important for us to know what sorts of numbers we can expect as output from a certain function. We call the set of all output values the *range* of the function.

Range

The set of all output values from a function. This is the set of all y -values or all possible values of the dependent variable.

The function $f(x) = |x|$ can accept any number as input, but only non-negative numbers ever come out. We say that the range of f is non-negative real numbers.

In the case of $g(x) = \frac{1}{x}$, notice that the only way for a fraction to equal 0 is if the numerator is equal to zero. No matter what nonzero input value we put into the function g , we will never get 0 as the output! The range of g is nonzero real numbers.

⁶ Did you know that the so-called “Hawaiian pizza”, a pizza topped with pieces of ham and pineapple, is not actually a Hawaiian invention? Several news organizations agree that this pizza originated in Ontario, Canada.

Piecewise definition

An interesting way to control the output of a function is to give it different definitions depending on the input. For example, suppose we define a function like so:

$$f(x) = \begin{cases} x & \text{if } x \text{ is positive} \\ -x & \text{if } x \text{ is negative} \\ 0 & \text{if } x \text{ is equal to } 0 \end{cases}$$

This unusual notation gives us several definitions at once, depending on which of the “if” statements applies to a given input value x .

If the input x is positive, then the output is simply x itself. If the input is negative, then the output is the opposite of x (so again, we get a positive value). Finally, in the third case, if the input is 0, then we get 0 as the output. Do you recognize this function? This case-by-case definition is really just a complicated way of saying $f(x) = |x|$.

Here’s a more interesting example. Suppose we define a function whose domain is the non-negative integers as follows:

$$z(n) = \begin{cases} \frac{n}{2} & \text{if } n \text{ is even} \\ -\frac{n+1}{2} & \text{if } n \text{ is odd} \end{cases}$$

This function produces different output depending on whether the input value is even or odd. Let’s try the first few input values to see what happens. What happens when we input 0? Zero is an even number, so the first case applies and:

$$z(0) = \frac{0}{2} = 0.$$

One is an odd number, so the second case applies:

$$z(1) = -\frac{1+1}{2} = -\frac{2}{2} = -1$$

Two is even, and three is odd. So we have:

$$z(2) = \frac{2}{2} = 1 \quad \text{and} \quad z(3) = -\frac{3+1}{2} = -\frac{4}{2} = -2.$$

This function produces the sequence:

$$0, -1, 1, -2, 2, -3, 3, -4, 4, \dots$$

which is a list of all the integers. In other words: this machine turns the natural numbers $\mathbb{N}$ (plus zero) into the set of integers $\mathbb{Z}$.⁷

⁷ This function suggests a shocking idea. Most people would say that there are more integers than there are natural numbers.

After all, since the integers include all of the natural numbers *and* all of the opposites of the natural numbers, it must be that the set of integers has *twice as many members*. . . right? Nope! The set of natural numbers and the set of integers are *exactly the same size*. Think of it this way: the function z defined above pairs up natural numbers and integers. If I give you a natural number (input), can you tell me what integer (output) it partners up with? If I give you an integer (output), can you tell me what natural number (input) it partners up with? If so, then the two sets match up exactly and neither set has any leftovers. They must therefore be the same size! Mind blown!

4.3 Three families of Algebra 1

The majority of Algebra 1 will focus on the in-depth study of three key families of functions: the linear family, the exponential family, and the quadratic family. We've encountered each of these families before, and now we will recap all that we've learned about them. Though this is our first "formal introduction", you might be surprised to discover how much you already know about the families!

4.3.1 Parent functions

For a given function to be considered a member of a certain family, it must share the characteristics of other members of that family as seen in its graph, its equation, and the patterns it exhibits in a table of values. These features are shown in their most basic form by the *parent function* of each family.

The linear parent is $y = x$ and the quadratic parent is $y = x^2$. In this course, we will consider $y = 2^x$ to be the parent function for the exponential family, though other choices are possible.⁸




Figure 4.2: The parent functions of the three families

Compare the equations for each of the parent functions: How are they the same? How are they different? Compare the graphs: How are they the same? How are they different? When presented with a new equation of graph, what features might we look for to decide which family it belongs to?⁹

In section 4.1 we learned techniques for identifying whether a given relation is a function. One of our tasks going forward will be to learn techniques for distinguishing the function families from one another.

⁸ A more natural choice for the exponential parent is the function $y = e^x$, where e is Euler's number (yes, that's the same Euler as the function notation guy). Euler's number has many important connections to the family of exponential functions, but those will have to wait until later. By the way e , like π , is an irrational number: $e \approx 2.71828\,18284\,59045\,23536\,02875\dots$

⁹ These are not the only three families of functions out there, of course. Other families of functions will become the subject of study in Algebra 2 (and beyond).

4.3.2 Transforming the parents

As we have seen many times already, a graph can be a particularly helpful way to see, and therefore understand, a mathematical relationship. One of the skills we'll develop as we go forward in this course is the ability to connect features of the equation of a certain function to features of its graph.

Over time and with practice, we'll get better at picturing the graph of a function in our mind's eye without having to draw the graph on paper. Drawing graphs will still be informative, of course, but some features of the graph will start to “jump out at us” simply from looking at the equation of the function.

Not far from the tree

// startup exploration

Study the three sets of animations depicted in figs. 4.3 to 4.5. Each set of images shows a way in which we can change the parent equation and produce a change in its graph.

Figure 4.3 shows what happens when we multiply the parent function by a constant value. Figure 4.4 shows what happens when we add a constant value to the parent. Figure 4.5 shows what happens when we add a constant value to the argument (the x -value) before that value is used as input to the parent function.

Study the animations^a or experiment with the interactive versions at algebranomicon.org. Then write a sentence or two summarizing what is happening in each case.

^a The images will only appear animated when this PDF document is viewed using Adobe Reader.

4.3.3 Fundamentals of linear functions

The first family we will study in depth is the family of linear functions. Here we summarize the key features of this family, all of which will return in future chapters.

We will learn several different forms for the equation of a linear function. The feature they all share is that the variable x has no exponent, or rather the highest power of the variable x is 1. For example, the parent of this family, $y = x$, is the same as $y = x^1$. That's a “phantom 1” there in the exponent, and we don't usually write it.

The graph of a linear function is a straight, non-vertical line. (Can you explain why we have to exclude vertical lines from this family of functions?) We sometimes also exclude horizontal lines from this family: a function of the form “ $y = \text{a number}$ ” is a horizontal line, for example $y = 6$.¹⁰

¹⁰ In fact, horizontal lines can be considered their own family: the family of *constant functions*.

Figure 4.3: Comparing the parent function $f(x)$ with offspring $a \cdot f(x)$.

Figure 4.4: Comparing the parent function $f(x)$ with offspring $f(x) + c$.

Figure 4.5: Comparing the parent function $f(x)$ with offspring $f(x + b)$.

Recall that arithmetic sequences are in the family of linear functions. Arithmetic sequences have a rule that involves repeatedly adding on a constant difference at each step. So, this is the feature we see in the table of a linear function: when the x -values increase by a fixed amount, the y -values increase by a (possibly different) fixed amount.

There are two key transformations of the parent function. Multiplying by a constant value, as in $y = a \cdot x$ changes the “steepness” of the graph. Certain a values cause the graph to flip over the x -axis and show a trend that goes in the opposite direction (we get a graph that has been reflected over the x -axis). Adding a value as in $y = x + c$ shifts the graph upwards or downwards. Viewed from a different perspective, we could also say that this shifts the graph left or right. Up-down shifts and left-right shifts look the same for a straight line. The distinction between these two transformations is more clear in the other families.

4.3.4 Fundamentals of exponential functions

The second family of functions we will explore in this course is the exponential family. In fact, we will begin our study of this family in Algebra 1 and then return to it again in Algebra 2.

An exponential equation is one in which the variable x appears as the exponent (hence the name). The graph of an exponential function is a smooth, curving shape that we sometimes describe as “J-like”. This is only an approximation, though, since an exponential graph may sometimes look more like a backwards J or an upside-down J.

Exponentials have the interesting property that they get almost vertical on one side and almost horizontal on the other side. For example, one side of the parent function, $y = 2^x$, gets “infinitely large” while on the other side we have tinier and tinier fractions that get closer and closer to zero. . . but never actually disappear! Mathematically speaking, we say that the function approaches an **asymptote** at the flat horizontal portion. The parent function has a horizontal asymptote at $y = 0$; other functions in this family may have their asymptote at a different position.

Recall that geometric sequences are in the family of exponential functions. Geometric sequences have a rule that involves repeatedly multiplying by a constant ratio. This is the feature we look for in a table.

Multiplying the parent by a value, as in $y = a \cdot 2^x$, changes the steepness of the graph and (if a is negative) causes it to be reflected over the x -axis.

Adding a value to the parent function, $y = 2^x + c$, causes a shift up or down. (Can you predict where the function’s asymptote will be in this case?) If we add to the argument, $y = 2^{(x+b)}$, we get a shift left or right.

4.3.5 Fundamentals of quadratic functions

We will study the quadratic family last in this course. In fact, it will be our gateway into the study of techniques that will become helpful as we learn about other families and types of functions in Algebra 2.

The key feature of a quadratic equation is that the highest exponent on the variable x is 2. The equation may also contain an x term, but that's the same as x^1 , so the highest power of x is still 2!

The graph of a quadratic function is a smooth U-like graph called a **parabola**. Of course, under certain circumstances, the U opens downwards.

The transformations of the parent function are perhaps more clear with this family. Multiplying by a constant, $y = a \cdot x^2$, changes the “steepness” of the graph. And, in the case that a is negative, controls whether the graph opens upwards or downwards.

Adding to the parent function $y = x^2 + c$ produces a clear vertical shift, and adding to the argument $y = (x + b)^2$ produces a clear horizontal shift. Interestingly, the horizontal shifts appear to be “backwards” compared to what we might expect. Adding to the argument shifts the graph to the left, whereas subtracting from the argument (or adding a negative value) shifts the graph to the right. (Can you spot this unusual behavior in the linear and exponential families, too?)

Recall from our study of sequences that quadratic patterns show a constant second difference. In other words, they have a recursive rule that states: “Start with A , add to each previous term the members of some arithmetic sequence.” This constant second difference will be the feature that we look for in a data table.

4.3.6 Domain and range of the three families

It's quite easy to identify the domains of the three main families of Algebra 1. Linear, exponential, and quadratic functions all have “all real numbers” as their domain. We can always use any number as input. This is convenient!

The range of a linear function is also pretty simple: it is all real numbers. As we saw as we transformed linear functions, the a (non-horizontal) line will eventually stretch up (or down) as high (or as low) as you want to go on the y -axis.

The ranges for exponential and quadratic functions are a little more complicated, but we can see what is happening by examining their graphs. Notice, for example, that the graphs of the parents $y = 2^x$ and $y = x^2$ never dip below the x -axis. This feature of the exponential and quadratic families will always be retained, in some form, when we transform these equations: the range will always be limited by some minimum or maximum value.

The quadratic parent passes through the origin, and so the range of $y = x^2$ is all real numbers *greater than or equal to* zero. We might write that y is any real number such that $y \geq 0$. In fact, every parabola has a point, called its **vertex**, that is either its lowest point (if the parabola opens upwards) or its highest point (if the parabola opens downwards). We will learn more about identifying and describing the vertex of a parabola when we study the quadratic family in in depth.

When it comes to the exponential parent, we mentioned that it has an asymptote at $y = 0$. This means that the function approaches, but never reaches, that value. So, the range of $y = 2^x$ is all real numbers *strictly greater than* zero. We might write that y is any real number such that $y > 0$. We will learn more about what's going on with this asymptote when we study the exponential family.

» Chapter summary «

In this chapter we have brought together a number of different concepts. Our work with sequences, graphs, functions, and transformations has been organized under three key headings: the three families that we will study closely in Algebra 1.

With the next chapter we begin to focus on the linear family, the most logical place to start. In one sense, the linear family is the easiest family to understand, and so it makes for a good “training level”. At the same time, the linear function is an important structural concept in mathematics that gives rise to a whole branch of study called, appropriately enough, *linear algebra*.

In any case, studying the linear family will give us the chance to learn a whole collection of new algebraic tools and techniques that we will use throughout the course. Onward!