

CHAPTER 1

Functions, Graphs, and Lines

Trying to do calculus without using functions would be one of the most pointless things you could do. If calculus had an ingredients list, functions would be first on it, and by some margin too. So, the first two chapters of this book are designed to jog your memory about the main features of functions. This chapter contains a review of the following topics:

- functions: their domain, codomain, and range, and the vertical line test;
- inverse functions and the horizontal line test;
- composition of functions;
- odd and even functions;
- graphs of linear functions and polynomials in general, as well as a brief survey of graphs of rational functions, exponentials, and logarithms; and
- how to deal with absolute values.

Trigonometric functions, or trig functions for short, are dealt with in the next chapter. So, let's kick off with a review of what a function actually is.

1.1 Functions

A *function* is a rule for transforming an object into another object. The object you start with is called the *input*, and comes from some set called the *domain*. What you get back is called the *output*; it comes from some set called the *codomain*.

Here are some examples of functions:

- Suppose you write $f(x) = x^2$. You have just defined a function f which transforms any number into its square. Since you didn't say what the domain or codomain are, it's assumed that they are both \mathbb{R} , the set of all real numbers. So you can square any real number, and get a real number back. For example, f transforms 2 into 4; it transforms $-1/2$ into $1/4$; and it transforms 1 into 1. This last one isn't much of a change at all, but that's no problem: the transformed object doesn't have to be different from the original one. When you write $f(2) = 4$, what you really mean

2 • Functions, Graphs, and Lines

is that f transforms 2 into 4. By the way, f is the transformation rule, while $f(x)$ is the result of applying the transformation rule to the variable x . So it's technically not correct to say " $f(x)$ is a function"; it should be " f is a function."

- Now, let $g(x) = x^2$ with domain consisting only of numbers greater than or equal to 0. (Such numbers are called *nonnegative*.) This seems like the same function as f , but it's not: the domains are different. For example, $f(-1/2) = 1/4$, but $g(-1/2)$ isn't defined. The function g just chokes on anything not in the domain, refusing even to touch it. Since g and f have the same rule, but the domain of g is smaller than the domain of f , we say that g is formed by *restricting the domain* of f .
- Still letting $f(x) = x^2$, what do you make of $f(\text{horse})$? Obviously this is undefined, since you can't square a horse. On the other hand, let's set

$$h(x) = \text{number of legs } x \text{ has,}$$

where the domain of h is the set of all animals. So $h(\text{horse}) = 4$, while $h(\text{ant}) = 6$ and $h(\text{salmon}) = 0$. The codomain could be the set of all nonnegative integers, since animals don't have negative or fractional numbers of legs. By the way, what is $h(2)$? This isn't defined, of course, since 2 isn't in the domain. How many legs does a "2" have, after all? The question doesn't really make sense. You might also think that $h(\text{chair}) = 4$, since most chairs have four legs, but that doesn't work either, since a chair isn't an animal, and so "chair" is not in the domain of h . That is, $h(\text{chair})$ is undefined.

- Suppose you have a dog called Junkster. Unfortunately, poor Junkster has indigestion. He eats something, then chews on it for a while and tries to digest it, fails, and hurls. Junkster has transformed the food into ... something else altogether. We could let

$$j(x) = \text{color of barf when Junkster eats } x,$$

where the domain of j is the set of foods that Junkster will eat. The codomain is the set of all colors. For this to work, we have to be confident that whenever Junkster eats a taco, his barf is always the same color (say, red). If it's sometimes red and sometimes green, that's no good: **a function must assign a unique output for each valid input.**

Now we have to look at the concept of the *range* of a function. The range is the set of all outputs that could possibly occur. You can think of the function working on transforming everything in the domain, one object at a time; the collection of transformed objects is the range. You might get duplicates, but that's OK.

So why isn't the range the same thing as the codomain? Well, the range is actually a subset of the codomain. The codomain is a set of **possible** outputs, while the range is the set of **actual** outputs. Here are the ranges of the functions we looked at above:

- If $f(x) = x^2$ with domain \mathbb{R} and codomain \mathbb{R} , the range is the set of nonnegative numbers. After all, when you square a number, the result cannot be negative. How do you know the range is **all** the nonnegative numbers? Well, if you square every number, you definitely cover all nonnegative numbers. For example, you get 2 by squaring $\sqrt{2}$ (or $-\sqrt{2}$).
- If $g(x) = x^2$, where the domain of g is only the nonnegative numbers but the codomain is still all of \mathbb{R} , the range will again be the set of nonnegative numbers. When you square every nonnegative number, you still cover all the nonnegative numbers.
- If $h(x)$ is the number of legs the animal x has, then the range is all the possible numbers of legs that **any** animal can have. I can think of animals that have 0, 2, 4, 6, and 8 legs, as well as some creepy-crawlies with more legs. If you include individual animals which have lost one or more legs, you can also include 1, 3, 5, and 7 in the mix, as well as other possibilities. In any case, the range of this function isn't so clear-cut; you probably have to be a biologist to know the real answer.
- Finally, if $j(x)$ is the color of Junkster's barf when he eats x , then the range consists of all possible barf-colors. I dread to think what these are, but probably bright blue isn't among them.

1.1.1 Interval notation

In the rest of this book, our functions will always have codomain \mathbb{R} , and the domain will always be as much of \mathbb{R} as possible (unless stated otherwise). So we'll often be dealing with subsets of the real line, especially connected intervals such as $\{x : 2 \leq x < 5\}$. It's a bit of a pain to write out the full set notation like this, but it sure beats having to say "all the numbers between 2 and 5, including 2 but not 5." We can do even better using interval notation.


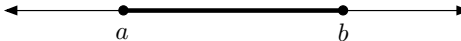







We'll write $[a, b]$ to mean the set of all numbers between a and b , including a and b themselves. So $[a, b]$ means the set of all x such that $a \leq x \leq b$. For example, $[2, 5]$ is the set of all real numbers between 2 and 5, including 2 and 5. (It's not just the set consisting of 2, 3, 4, and 5: don't forget that there are loads of fractions and irrational numbers between 2 and 5, such as $5/2$, $\sqrt{7}$, and π .) An interval such as $[a, b]$ is called *closed*.

If you don't want the endpoints, change the square brackets to parentheses. In particular, (a, b) is the set of all numbers between a and b , not including a or b . So if x is in the interval (a, b) , we know that $a < x < b$. The set $(2, 5)$ includes all real numbers between 2 and 5, but not 2 or 5. An interval of the form (a, b) is called *open*.

You can mix and match: $[a, b)$ consists of all numbers between a and b , including a but not b . And $(a, b]$ includes b but not a . These intervals are closed at one end and open at the other. Sometimes such intervals are called *half-open*. An example is the set $\{x : 2 \leq x < 5\}$ from above, which can also be written as $[2, 5)$.

There's also the useful notation (a, ∞) for all the numbers greater than a not including a ; $[a, \infty)$ is the same thing but with a included. There are three other possibilities which involve $-\infty$; all in all, the situation looks like this:

4 • Functions, Graphs, and Lines

(a, b)	$\{x : a < x < b\}$	
$[a, b]$	$\{x : a \leq x \leq b\}$	
$(a, b]$	$\{x : a < x \leq b\}$	
$[a, b)$	$\{x : a \leq x < b\}$	
(a, ∞)	$\{x : x > a\}$	
$[a, \infty)$	$\{x : x \geq a\}$	
$(-\infty, b)$	$\{x : x < b\}$	
$(-\infty, b]$	$\{x : x \leq b\}$	
$(-\infty, \infty)$	\mathbb{R}	

1.1.2 Finding the domain

Sometimes the definition of a function will include the domain. (This was the case, for example, with our function g from Section 1.1 above.) Most of the time, however, the domain is not provided. The basic convention is that the domain consists of as much of the set of real numbers as possible. For example, if $k(x) = \sqrt{x}$, the domain can't be all of \mathbb{R} , since you can't take the square root of a negative number. The domain must be $[0, \infty)$, which is just the set of all numbers greater than or equal to 0.



OK, so square roots of negative numbers are bad. What else can cause a screw-up? Here's a list of the three most common possibilities:

1. The denominator of a fraction can't be zero.
2. You can't take the square root (or fourth root, sixth root, and so on) of a negative number.
3. You can't take the logarithm of a negative number or of 0. (Remember logs? If not, see Chapter 9!)

You might recall that $\tan(90^\circ)$ is also a problem, but this is really a special case of the first item above. You see,

$$\tan(90^\circ) = \frac{\sin(90^\circ)}{\cos(90^\circ)} = \frac{1}{0},$$



so the reason $\tan(90^\circ)$ is undefined is really that a hidden denominator is zero. Here's another example: if we try to define

$$f(x) = \frac{\log_{10}(x+8)\sqrt{26-2x}}{(x-2)(x+19)},$$

then what is the domain of f ? Well, for $f(x)$ to make sense, here's what needs to happen:

- We need to take the square root of $(26 - 2x)$, so this quantity had better be nonnegative. That is, $26 - 2x \geq 0$. This can be rewritten as $x \leq 13$.

Section 1.1.3: Finding the range using the graph • 5


- We also need to take the logarithm of $(x + 8)$, so this quantity needs to be positive. (Notice the difference between logs and square roots: you can take the square root of 0, but you can't take the log of 0.) Anyway, we need $x + 8 > 0$, so $x > -8$. So far, we know that $-8 < x \leq 13$, so the domain is at most $(-8, 13]$.
- The denominator can't be 0; this means that $(x - 2) \neq 0$ and $(x + 19) \neq 0$. In other words, $x \neq 2$ and $x \neq -19$. This last one isn't a problem, since we already know that x lies in $(-8, 13]$, so x can't possibly be -19 . We do have to exclude 2, though.

So we have found that the domain is the set $(-8, 13]$ except for the number 2. This set could be written as $(-8, 13] \setminus \{2\}$. Here the backslash means "not including."

1.1.3 Finding the range using the graph

Let's define a new function F by specifying that its domain is $[-2, 1]$ and that $F(x) = x^2$ on this domain. (Remember, the codomain of any function we look at will always be the set of all real numbers.) Is F the same function as f , where $f(x) = x^2$ for all real numbers x ? The answer is no, since the two functions have different domains (even though they have the same rule). As in the case of the function g from Section 1.1 above, the function F is formed by restricting the domain of f .

Now, what is the range of F ? Well, what happens if you square every number between -2 and 1 inclusive? You should be able to work this out directly, but this is a good opportunity to see how to use a graph to find the range of a function. The idea is to sketch the graph of the function, then imagine two rows of lights shining from the far left and far right of the graph horizontally toward the y -axis. The curve will cast two shadows, one on the left side and one on the right side of the y -axis. The range is the union of both shadows: that is, if any point on the y -axis lies in either the left-hand or the right-hand shadow, it is in the range of the function. Let's see how this works with our function F :



6 • Functions, Graphs, and Lines


The left-hand shadow covers all the points on the y -axis between 0 and 4 inclusive, which is $[0, 4]$; on the other hand, the right-hand shadow covers the points between 0 and 1 inclusive, which is $[0, 1]$. The right-hand shadow doesn't contribute anything extra: the total coverage is still $[0, 4]$. This is the range of F .

1.1.4 The vertical line test

In the last section, we used the graph of a function to find its range. The graph of a function is very important: it really shows you what the function “looks like.” We'll be looking at techniques for sketching graphs in Chapter 12, but for now I'd like to remind you about the vertical line test.

You can draw any figure you like on a coordinate plane, but the result may not be the graph of a function. So what's special about the graph of a function? What is the graph of a function f , anyway? Well, it's the collection of all points with coordinates $(x, f(x))$, where x is in the domain of f . Here's another way of looking at this: start with some number x . If x is in the domain, you plot the point $(x, f(x))$, which of course is at a height of $f(x)$ units above the point x on the x -axis. If x isn't in the domain, you don't plot anything. Now repeat for every real number x to build up the graph.


Here's the key idea: you can't have two points with the same x -coordinate. In other words, no two points on the graph can lie on the same vertical line. Otherwise, how would you know which of the two or more heights above the point x on the x -axis corresponds to the value of $f(x)$? So, this leads us to the *vertical line test*: if you have some graph and you want to know whether it's the graph of a function, see whether any vertical line intersects the graph more than once. If so, it's not the graph of a function; but if no vertical line intersects the graph more than once, you are indeed dealing with the graph of a function. For example, the circle of radius 3 units centered at the origin has a graph like this:



Such a commonplace object should be a function, right? No, check the vertical lines that are shown in the diagram. Sure, to the left of -3 or to the right of 3 , there's no problem—the vertical lines don't even hit the graph, which is fine. Even at -3 or 3 , the vertical lines only intersect the curve in one point each, which is also fine. The problem is when x is in the interval $(-3, 3)$. For

any of these values of x , the vertical line through $(x, 0)$ intersects the circle twice, which screws up the circle's potential function-status. You just don't know whether $f(x)$ is the top point or the bottom point.

The best way to salvage the situation is to chop the circle in half horizontally and choose only the top or the bottom half. The equation for the whole circle is $x^2 + y^2 = 9$, whereas the equation for the top semicircle is $y = \sqrt{9 - x^2}$. The bottom semicircle has equation $y = -\sqrt{9 - x^2}$. These last two are functions, both with domain $[-3, 3]$. If you felt like chopping in a different way, you wouldn't actually have to take semicircles—you could chop and change between the upper and lower semicircles, as long as you don't violate the vertical line test. For example, here's the graph of a function which also has domain $[-3, 3]$:



The vertical line test checks out, so this is indeed the graph of a function.

1.2 Inverse Functions

Let's say you have a function f . You present it with an input x ; provided that x is in the domain of f , you get back an output, which we call $f(x)$. Now we try to do things all backward and ask this question: if you pick a number y , what input can you give to f in order to get back y as your output?

Here's how to state the problem in math-speak: given a number y , what x in the domain of f satisfies $f(x) = y$? The first thing to notice is that y has to be in the range of f . Otherwise, by definition there are no values of x such that $f(x) = y$. There would be nothing in the domain that f would transform into y , since the range is **all** the possible outputs.

On the other hand, if y is in the range, there might be many values that work. For example, if $f(x) = x^2$ (with domain \mathbb{R}), and we ask what value of x transforms into 64, there are obviously two values of x : 8 and -8 . On the other hand, if $g(x) = x^3$, and we ask the same question, there's only one value of x , which is 4. The same would be true for any number we give to g to transform, because any number has only one (real) cube root.

So, here's the situation: we're given a function f , and we pick y in the range of f . Ideally, there will be exactly one value of x which satisfies $f(x) = y$. If this is true for every value of y in the range, then we can define a new

8 • Functions, Graphs, and Lines

function which reverses the transformation. Starting with the output y , the new function finds the one and only input x which leads to the output. The new function is called the *inverse function of f* , and is written as f^{-1} . Here's a summary of the situation in mathematical language:

1. Start with a function f such that for any y in the range of f , there is exactly one number x such that $f(x) = y$. That is, different inputs give different outputs. Now we will define the inverse function f^{-1} .
2. The domain of f^{-1} is the same as the range of f .
3. The range of f^{-1} is the same as the domain of f .
4. The value of $f^{-1}(y)$ is the number x such that $f(x) = y$. So,

$$\text{if } f(x) = y, \quad \text{then } f^{-1}(y) = x.$$


The transformation f^{-1} acts like an undo button for f : if you start with x and transform it into y using the function f , then you can undo the effect of the transformation by using the inverse function f^{-1} on y to get x back.

This raises some questions: how do you see if there's only one value of x that satisfies the equation $f(x) = y$? If so, how do you find the inverse, and what does its graph look like? If not, how do you salvage the situation? We'll answer these questions in the next three sections.

1.2.1 The horizontal line test

For the first question—how to see that there's only one value of x that works for any y in the range—perhaps the best way is to look at the graph of your function. We want to pick y in the range of f and hopefully only have one value of x such that $f(x) = y$. What this means is that the horizontal line through the point $(0, y)$ should intersect the graph exactly once, at some point (x, y) . That x is the one we want. If the horizontal line intersects the curve more than once, there would be multiple potential inverses x , which is bad. In that case, the only way to get an inverse function is to restrict the domain; we'll come back to this very shortly. What if the horizontal line doesn't intersect the curve at all? Then y isn't in the range after all, which is OK.


So, we have just described the *horizontal line test*: if every horizontal line intersects the graph of a function at most once, the function has an inverse. If even one horizontal line intersects the graph more than once, there isn't an inverse function. For example, look at the graphs of $f(x) = x^3$ and $g(x) = x^2$:



No horizontal line hits $y = f(x)$ more than once, so f has an inverse. On the other hand, some of the horizontal lines hit the curve $y = g(x)$ twice, so g has no inverse. Here's the problem: if you want to solve $y = x^2$ for x , where y is positive, then there are two solutions, $x = \sqrt{y}$ and $x = -\sqrt{y}$. You don't know which one to take!

1.2.2 Finding the inverse

Now let's move on to the second of our questions: how do you find the inverse of a function f ? Well, you write down $y = f(x)$ and try to solve for x . In our example of $f(x) = x^3$, we have $y = x^3$, so $x = \sqrt[3]{y}$. This means that $f^{-1}(y) = \sqrt[3]{y}$. If the variable y here offends you, by all means switch it to x : you can write $f^{-1}(x) = \sqrt[3]{x}$ if you prefer. Of course, solving for x is not always easy and in fact is often impossible. On the other hand, if you know what the graph of your function looks like, the graph of the inverse function is easy to find. The idea is to draw the line $y = x$ on the graph, then pretend that this line is a two-sided mirror. The inverse function is the reflection of the original function in this mirror. When $f(x) = x^3$, here's what f^{-1} looks like:




The original function f is reflected in the mirror $y = x$ to get the inverse function. Note that the domain and range of both f and f^{-1} are the whole real line.


1.2.3 Restricting the domain

Finally, we'll address our third question: if the horizontal line test fails and there's no inverse, what can be done? Our problem is that there are multiple values of x that give the same y . The only way to get around the problem is to throw away all but one of these values of x . That is, we have to decide which one of our values of x we want to keep, and throw the rest away. As we saw in Section 1.1 above, this is called *restricting the domain* of our function. Effectively, we ghost out part of the curve so that what's left no longer fails the horizontal line test. For example, if $g(x) = x^2$, we can ghost out the left half of the graph like this:

10 • Functions, Graphs, and Lines




The new (unghosted) curve has the reduced domain $[0, \infty)$ and satisfies the horizontal line test, so there is an inverse function. More precisely, the function h , which has domain $[0, \infty)$ and is defined by $h(x) = x^2$ on this domain, has an inverse. Let's play the reflection game to see what it looks like:



To find the equation of the inverse, we have to solve for x in the equation $y = x^2$. Clearly the solution is $x = \sqrt{y}$ or $x = -\sqrt{y}$, but which one do we need? We know that the range of the inverse function is the same as the domain of the original function, which we have restricted to be $[0, \infty)$. So we need a nonnegative number as our answer, and that has to be $x = \sqrt{y}$. That is, $h^{-1}(y) = \sqrt{y}$. Of course, we could have ghosted out the right half of the original graph to restrict the domain to $(-\infty, 0]$. In that case, we'd get a function j which has domain $(-\infty, 0]$ and again satisfies $j(x) = x^2$, but only on this domain. This function also has an inverse, but the inverse is now the negative square root: $j^{-1}(y) = -\sqrt{y}$.

By the way, if you take the original function g given by $g(x) = x^2$ with domain $(-\infty, \infty)$, which fails the horizontal line test, and try to reflect it in the mirror $y = x$, you get the following picture:



Notice that the graph fails the vertical line test, so it's not the graph of a function. This illustrates the connection between the vertical and horizontal line tests—when horizontal lines are reflected in the mirror $y = x$, they become vertical lines.

1.2.4 Inverses of inverse functions

One more thing about inverse functions: if f has an inverse, it's true that $f^{-1}(f(x)) = x$ for all x in the domain of f , and also that $f(f^{-1}(y)) = y$ for all y in the range of f . (Remember, the range of f is the same as the domain of f^{-1} , so you can indeed take $f^{-1}(y)$ for y in the range of f without causing any screwups.)



For example, if $f(x) = x^3$, then f has an inverse given by $f^{-1}(x) = \sqrt[3]{x}$, and so $f^{-1}(f(x)) = \sqrt[3]{x^3} = x$ for any x . Remember, the inverse function is like an undo button. We use x as an input to f , and then give the output to f^{-1} ; this undoes the transformation and gives us back x , the original number. Similarly, $f(f^{-1}(y)) = (\sqrt[3]{y})^3 = y$. So f^{-1} is the inverse function of f , and f is the inverse function of f^{-1} . In other words, the inverse of the inverse is the original function.

Now, you have to be careful in the case where you restrict the domain. Let $g(x) = x^2$; we've seen that you need to restrict the domain to get an inverse. Let's say we restrict the domain to $[0, \infty)$ and carelessly continue to refer to the function as g instead of h , as in the previous section. We would then say that $g^{-1}(x) = \sqrt{x}$. If you calculate $g(g^{-1}(x))$, you find that this is $(\sqrt{x})^2$, which equals x , provided that $x \geq 0$. (Otherwise you can't take the square root in the first place.)

On the other hand, if you work out $g^{-1}(g(x))$, you get $\sqrt{x^2}$, which is not always the same thing as x . For example, if $x = -2$, then $x^2 = 4$ and so $\sqrt{x^2} = \sqrt{4} = 2$. So it's not true in general that $g^{-1}(g(x)) = x$. The problem is that -2 isn't in the restricted-domain version of g . Technically, you can't even compute $g(-2)$, since -2 is no longer in the domain of g . We really should be working with h , not g , so that we remember to be more careful. Nevertheless, in practice, mathematicians will often restrict the domain without changing letters! So it will be useful to summarize the situation as follows:



If the domain of a function f can be restricted so that f has an inverse f^{-1} , then

- $f(f^{-1}(y)) = y$ for all y in the range of f ; but
- $f^{-1}(f(x))$ may not equal x ; in fact, $f^{-1}(f(x)) = x$ only when x is in the restricted domain.

We'll be revisiting these important points in the context of inverse trig functions in Section 10.2.6 of Chapter 10.

1.3 Composition of Functions

Let's say we have a function g given by $g(x) = x^2$. You can replace x by anything you like, as long as it makes sense. For example, you can write

12 • Functions, Graphs, and Lines

$g(y) = y^2$, or $g(x + 5) = (x + 5)^2$. This last example shows that you need to be very careful with parentheses. It would be wrong to write $g(x+5) = x+5^2$, since this is just $x + 25$, which is not the same thing as $(x + 5)^2$. If in doubt, use parentheses. That is, if you need to write out $f(\text{something})$, replace every instance of x by (something) , making sure to include the parentheses. Just about the only time you don't need to use parentheses is when the function is an exponential function—for example, if $h(x) = 3^x$, then you can just write $h(x^2 + 6) = 3^{x^2+6}$. You don't need parentheses since you're already writing the $x^2 + 6$ as a superscript.

Now consider the function f defined by $f(x) = \cos(x^2)$. If I give you a number x , how do you compute $f(x)$? Well, first you square it, then you take the cosine of the result. Since we can decompose the action of $f(x)$ into these two separate actions which are performed one after the other, we might as well describe those actions as functions themselves. So, let $g(x) = x^2$ and $h(x) = \cos(x)$. To simulate what f does when you use x as an input, you could first give x to g to square it, and then instead of taking the result back you could ask g to give its result to h instead. Then h spits out a number, which is the final answer. The answer will, of course, be the cosine of what came out of g , which was the square of the original x . This behavior exactly mimics f , so we can write $f(x) = h(g(x))$. Another way of expressing this is to write $f = h \circ g$; here the circle means “composed with.” That is, f is h composed with g , or in other words, f is the *composition* of h and g . What's tricky is that you write h before g (reading from left to right as usual!) but you apply g first. I agree that it's confusing, but what can I say—you just have to deal with it.



It's useful to practice composing two or more functions together. For example, if $g(x) = 2^x$, $h(x) = 5x^4$, and $j(x) = 2x - 1$, what is a formula for the function $f = g \circ h \circ j$? Well, just replace one thing at a time, starting with j , then h , then g . So:

$$f(x) = g(h(j(x))) = g(h(2x - 1)) = g(5(2x - 1)^4) = 2^{5(2x-1)^4}.$$



You should also practice reversing the process. For example, suppose you start off with

$$f(x) = \frac{1}{\tan(5 \log_2(x + 3))}.$$

How would you decompose f into simpler functions? Zoom in to where you see the quantity x . The first thing you do is add 3, so let $g(x) = x + 3$. Then you have to take the base 2 logarithm of the resulting quantity, so set $h(x) = \log_2(x)$. Next, multiply by 5, so set $j(x) = 5x$. Then take the tangent, so put $k(x) = \tan(x)$. Finally, take reciprocals, so let $m(x) = 1/x$. With all these definitions, you should check that

$$f(x) = m(k(j(h(g(x)))).$$

Using the composition notation, you can write

$$f = m \circ k \circ j \circ h \circ g.$$

This isn't the only way to break down f . For example, we could have combined h and j into another function n , where $n(x) = 5 \log_2(x)$. Then you should check that $n = j \circ h$, and

$$f = m \circ k \circ n \circ g.$$

Perhaps the original decomposition (involving j and h) is better because it breaks down f into more elementary steps, but the second one (involving n) isn't wrong. After all, $n(x) = 5 \log_2(x)$ is still a pretty simple function of x .

Beware: composition of functions isn't the same thing as multiplying them together. For example, if $f(x) = x^2 \sin(x)$, then f is not the composition of two functions. To calculate $f(x)$ for any given x , you actually have to find both x^2 and $\sin(x)$ (it doesn't matter which one you find first, unlike with composition) and then multiply these two things together. If $g(x) = x^2$ and $h(x) = \sin(x)$, then we'd write $f(x) = g(x)h(x)$, or $f = gh$. Compare this to the composition of the two functions, $j = g \circ h$, which is given by


$$j(x) = g(h(x)) = g(\sin(x)) = (\sin(x))^2$$

or simply $j(x) = \sin^2(x)$. The function j is a completely different function from the product $x^2 \sin(x)$. It's also different from the function $k = h \circ g$, which is also a composition of g and h but in the other order:

$$k(x) = h(g(x)) = h(x^2) = \sin(x^2).$$

This is yet another completely different function. The moral of the story is that products and compositions are not the same thing, and furthermore, the order of the functions matters when you compose them, but not when you multiply them together.

One simple but important example of composition of functions occurs when you compose some function f with $g(x) = x - a$, where a is some constant number. You end up with a new function h given by $h(x) = f(x - a)$. A useful point to note is that the graph of $y = h(x)$ is the same as the graph of $y = f(x)$, except that it's shifted over a units to the right. If a is negative, then the shift is to the left. (The way to think of this, for example, is that a shift of -3 units to the right is the same as a shift of 3 units to the left.) So, how would you sketch the graph of $y = (x - 1)^2$? This is the same as $y = x^2$, but with x replaced by $x - 1$. So the graph of $y = x^2$ needs to be shifted to the right by 1 unit, and looks like this:



1.4 • Functions, Graphs, and Lines

Similarly, the graph of $y = (x + 2)^2$ is the graph of $y = x^2$ shifted to the left by 2 units, since you can interpret $(x + 2)$ as $(x - (-2))$.

1.4 Odd and Even Functions


Some functions have some symmetry properties that make them easier to deal with. Consider the function f given by $f(x) = x^2$. Pick any positive number you like (I'll choose 3) and hit it with f (I get 9). Now take the negative of that number, -3 in my case, and hit that with f (I get 9 again). You should get the same answer both times, as I did, regardless of which number you chose. You can express this phenomenon by writing $f(-x) = f(x)$ for all x . That is, if you give x to f as an input, you get back the same answer as if you used the input $-x$ instead. Notice that $g(x) = x^4$ and $h(x) = x^6$ also have this property—in fact, $j(x) = x^n$, where n is any even number (n could in fact be negative), has the same property. Inspired by this, we say that a function f is *even* if $f(-x) = f(x)$ for all x in the domain of f . It's not good enough for this equation to be true for some values of x ; it has to be true for **all** x in the domain of f .

Now, let's say we play the same game with $f(x) = x^3$. Take your favorite positive number (I'll stick with 3) and hit that with f (I get 27). Now try again with the negative of your number, -3 in my case; I get -27 , and you should also get the negative of what you got before. You can express this mathematically as $f(-x) = -f(x)$. Once again, the same property holds for $j(x) = x^n$ when n is any odd number (and once again, n could be negative). So, we say that a function f is *odd* if $f(-x) = -f(x)$ for all x in the domain of f .




In general, a function might be odd, it might be even, or it might be neither odd nor even. Don't forget this last point! Most functions are neither odd nor even. On the other hand, there's only one function that's both odd and even, which is the rather boring function given by $f(x) = 0$ for all x (we'll call this the "zero function"). Why is this the only odd and even function? Let's convince ourselves. If the function f is even, then $f(-x) = f(x)$ for all x . But if it's also odd, then $f(-x) = -f(x)$ for all x . Take the first of these equations and subtract the second from it. You should get $0 = 2f(x)$, which means that $f(x) = 0$. This is true for all x , so the function f must just be the zero function. One other nice observation is that if a function f is odd, and the number 0 is in its domain, then $f(0) = 0$. Why is it so? Because $f(-x) = -f(x)$ is true for all x in the domain of f , so let's try it for $x = 0$. You get $f(-0) = -f(0)$. But -0 is the same thing as 0, so we have $f(0) = -f(0)$. This simplifies to $2f(0) = 0$, or $f(0) = 0$ as claimed.

Anyway, starting with a function f , how can you tell if it is odd, even, or neither? And so what if it is odd or even anyway? Let's look at this second question before coming back to the first one. One nice thing about knowing that a function is odd or even is that it's easier to graph the function. In fact, if you can graph the right-hand half of the function, the left-hand half is a piece of cake! Let's say that f is an even function. Then since $f(x) = f(-x)$, the graph of $y = f(x)$ is at the same height above the x -coordinates x and $-x$. This is true for all x , so the situation looks something like this:



We can conclude that **the graph of an even function has mirror symmetry about the y -axis**. So, if you graph the right half of a function which you know is even, you can get the left half by reflecting the right half about the y -axis. Check the graph of $y = x^2$ to make sure that it has this mirror symmetry.

On the other hand, let's say that f is an odd function. Since we have $f(-x) = -f(x)$, the graph of $y = f(x)$ is at the same height **above** the x -coordinate x as it is **below** the x -coordinate $-x$. (Of course, if $f(x)$ is negative, then you have to switch the words "above" and "below.") In any case, the picture looks like this:



The symmetry is now a point symmetry about the origin. That is, **the graph of an odd function has 180° point symmetry about the origin**. This means that if you only have the right half of a function which you know is odd, you can get the left half as follows. Pretend that the curve is sitting on top of the paper, so you can pick it up if you like but you can't change its shape. Instead of picking it up, put a pin through the curve at the origin (remember, odd functions must pass through the origin if they are defined at 0) and then spin the whole curve around half a revolution. This is what the left-hand half of the graph looks like. (This doesn't work so well if the curve isn't continuous, that is, if the curve isn't all in one piece!) Check to see that the above graph and also the graph of $y = x^3$ have this symmetry.



Now, suppose f is defined by the equation $f(x) = \log_5(2x^6 - 6x^2 + 3)$. How do you tell if f is odd, even, or neither? The technique is to calculate $f(-x)$ by replacing every instance of x with $(-x)$, making sure not to forget the parentheses around $-x$, and then simplifying the result. If you end up with the original expression $f(x)$, then f is even; if you end up with the negative of the original expression $-f(x)$, then f is odd; if you end up with a mess that isn't either $f(x)$ or $-f(x)$, then f is neither (or you didn't simplify enough!).

In the example above, you'd write

$$f(-x) = \log_5(2(-x)^6 - 6(-x)^2 + 3) = \log_5(2x^6 - 6x^2 + 3),$$



which is actually equal to the original $f(x)$. So the function f is even. How about

$$g(x) = \frac{2x^3 + x}{3x^2 + 5} \quad \text{and} \quad h(x) = \frac{2x^3 + x - 1}{3x^2 + 5}?$$

Well, for g , we have

$$g(-x) = \frac{2(-x)^3 + (-x)}{3(-x)^2 + 5} = \frac{-2x^3 - x}{3x^2 + 5}.$$

Now you have to observe that you can take the minus sign out front and write

$$g(-x) = -\frac{2x^3 + x}{3x^2 + 5},$$

which, you notice, is equal to $-g(x)$. That is, apart from the minus sign, we get the original function back. So, g is an odd function. How about h ? We have

$$h(-x) = \frac{2(-x)^3 + (-x) - 1}{3(-x)^2 + 5} = \frac{-2x^3 - x - 1}{3x^2 + 5}.$$

Once again, we take out the minus sign to get

$$h(-x) = -\frac{2x^3 + x + 1}{3x^2 + 5}.$$

Hmm, this doesn't appear to be the negative of the original function, because of the $+1$ term in the numerator. It's not the original function either, so the function h is neither odd nor even.



Let's look at one more example. Suppose you want to prove that the product of two odd functions is always an even function. How would you go about doing this? Well, it helps to have names for things, so let's say we have two odd functions f and g . We need to look at the product of these functions, so let's call the product h . That is, we define $h(x) = f(x)g(x)$. So, our task is to show that h is even. We'll do this by showing that $h(-x) = h(x)$, as usual. It will be helpful to note that $f(-x) = -f(x)$ and $g(-x) = -g(x)$, since f and g are odd. Let's start with $h(-x)$. Since h is the product of f and g , we have $h(-x) = f(-x)g(-x)$. Now we use the oddness of f and g to express this last term as $(-f(x))(-g(x))$. The minus signs now come out front and cancel out, so this is the same thing as $f(x)g(x)$ which of course equals $h(x)$. We could (and should) express all this text mathematically like this:


$$h(-x) = f(-x)g(-x) = (-f(x))(-g(x)) = f(x)g(x) = h(x).$$



Anyway, since $h(-x) = h(x)$, the function h is even. Now you should try to prove that the product of two even functions is always even, and also that the product of an odd and an even function must be odd. Go on, do it now!


1.5 Graphs of Linear Functions

Functions of the form $f(x) = mx + b$ are called *linear*. There's a good reason for this: the graphs of these functions are lines. (As far as we're concerned, the word "line" always means "straight line.") The slope of the line is given by m . Imagine for a moment that you are in the page, climbing the line as if it were a mountain. You start at the left side of the page and head to the right, like this:



If the slope m is positive, as it is in the above picture, then you are heading uphill. The bigger m is, the steeper the climb. On the other hand, if the slope is negative, then you are heading downhill. The more negative the slope, the steeper the downhill grade. If the slope is zero, then the line is flat, or horizontal—you're going neither uphill nor downhill, just trudging along a flat line.


To sketch the graph of a linear function, you only need to identify two points on the graph. This is because there's only one line that goes through two different points. You just put your ruler on the points and draw the line. One point is easy to find, namely, the y -intercept. Set $x = 0$ in the equation $y = mx + b$, and you see that $y = m \times 0 + b = b$. That is, the y -intercept is equal to b , so the line goes through $(0, b)$. To find another point, you could find the x -intercept by setting $y = 0$ and finding what x is. This works pretty well except in two cases. The first case is when $b = 0$, in which case we are just dealing with $y = mx$. This goes through the origin, so the x -intercept and the y -intercept are both zero. To get another point, you'll just have to substitute in $x = 1$ and see that $y = m$. So, the line $y = mx$ goes through the origin and $(1, m)$. For example, the line $y = -2x$ goes through the origin and also through $(1, -2)$, so it looks like this:



The other bad case is when $m = 0$. But then we just have $y = b$, which is a horizontal line through $(0, b)$.



For a more interesting example, consider $y = \frac{1}{2}x - 1$. The y -intercept is -1 , and the slope is $\frac{1}{2}$. To sketch the line, find the x -intercept by setting $y = 0$. We get $0 = \frac{1}{2}x - 1$, which simplifies to $x = 2$. So, the line looks like this:



Now, let's suppose you know that you have a line in the plane, but you don't know its equation. If you know it goes through a certain point, and you know what its slope is, then you can find the equation of the line. You really, really, really need to know how to do this, since it comes up a lot. This formula, called the *point-slope* form of a linear function, is what you need to know:

If a line goes through (x_0, y_0) and has slope m , then its equation is $y - y_0 = m(x - x_0)$.



For example, what is the equation of the line through $(-2, 5)$ which has slope -3 ? It is $y - 5 = -3(x - (-2))$, which you can expand and simplify down to $y = -3x - 1$.

Sometimes you don't know the slope of the line, but you do know two points that it goes through. How do you find the equation? The technique is to find the slope, then use the previous idea with one of the points (your choice) to find the equation. First, you need to know this:

If a line goes through (x_1, y_1) and (x_2, y_2) , its slope is equal to $\frac{y_2 - y_1}{x_2 - x_1}$.



So, what is the equation of the line through $(-3, 4)$ and $(2, -6)$? Let's find the slope first:

$$\text{slope} = \frac{-6 - 4}{2 - (-3)} = \frac{-10}{5} = -2.$$

We now know that the line goes through $(-3, 4)$ and has slope -2 , so its equation is $y - 4 = -2(x - (-3))$, or after simplifying, $y = -2x - 2$. Alternatively, we could have used the other point $(2, -6)$ with slope -2 to see that the equation of the line is $y - (-6) = -2(x - 2)$, which simplifies to $y = -2x - 2$. Thankfully this is the same equation as before—it doesn't matter which point you pick, as long as you have used both points to find the slope.

1.6 Common Functions and Graphs


Here are the most important functions you should know about.

1. Polynomials: these are functions built out of nonnegative integer powers of x . You start with the building blocks 1 , x , x^2 , x^3 , and so on, and you are allowed to multiply these basic functions by numbers and add a finite number of them together. For example, the polynomial $f(x) = 5x^4 - 4x^3 + 10$ is formed by taking 5 times the building block x^4 , and -4 times the building block x^3 , and 10 times the building block 1 , and adding them together. You might also want to include the intermediate building blocks x^2 and x , but since they don't appear, you need to take 0 times of each. The amount that you multiply the building block x^n by is called the *coefficient* of x^n . For example, in the polynomial f above, the coefficient of x^4 is 5, the coefficient of x^3 is -4 , the coefficients of x^2 and x are both 0, and the coefficient of 1 is 10. (Why allow x and 1, by the way? They seem different from the other blocks, but they're not really: $x = x^1$ and $1 = x^0$.) The highest number n such that x^n has a nonzero coefficient is called the *degree* of the polynomial. For example, the degree of the above polynomial f is 4, since no power of x greater than 4 is present. The mathematical way to write a general polynomial of degree n is

$$p(x) = a_n x^n + a_{n-1} x^{n-1} + \cdots + a_2 x^2 + a_1 x + a_0,$$


where a_n is the coefficient of x^n , a_{n-1} is the coefficient of x^{n-1} , and so on down to a_0 , which is the coefficient of 1.

Since the functions x^n are the building blocks of all polynomials, you should know what their graphs look like. The even powers mostly look similar to each other, and the same can be said for the odd powers. Here's what the graphs look like, from x^0 up to x^7 :



20 • Functions, Graphs, and Lines

Sketching the graphs of more general polynomials is more difficult. Even finding the x -intercepts is often impossible unless the polynomial is very simple. There is one aspect of the graph that is fairly straightforward, which is what happens at the far left and right sides of the graph. This is determined by the so-called *leading coefficient*, which is the coefficient of the highest-degree term. This is basically the number a_n defined above. For example, in our polynomial $f(x) = 5x^4 - 4x^3 + 10$ from above, the leading coefficient is 5. In fact, it only matters whether the leading coefficient is positive or negative. It also matters whether the degree of the polynomial is odd or even; so there are four possibilities for what the edges of the graph can look like:



The wiggles in the center of these diagrams aren't relevant—they depend on the other terms of the polynomial. The diagram is just supposed to show what the graphs look like near the left and right edges. In this sense, the graph of our polynomial $f(x) = 5x^4 - 4x^3 + 10$ looks like the leftmost picture above, since $n = 4$ is even and $a_n = 5$ is positive.

Let's spend a little time on degree 2 polynomials, which are called *quadratics*. Instead of writing $p(x) = a_2x^2 + a_1x + a_0$, it's easier to write the coefficients as a , b , and c , so we have $p(x) = ax^2 + bx + c$. Quadratics have two, one, or zero (real) roots, depending on the sign of the *discriminant*. The discriminant, which is often written as Δ , is given by $\Delta = b^2 - 4ac$. There are three possibilities. If $\Delta > 0$, then there are two roots; if $\Delta = 0$, there is one root, which is called a *double root*; and if $\Delta < 0$, then there are no roots. In the first two cases, the roots are given by

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$



Notice that the expression in the square root is just the discriminant. An important technique for dealing with quadratics is *completing the square*. Here's how it works. We'll use the example of the quadratic $2x^2 - 3x + 10$. The first step is to take out the leading coefficient as a factor. So our quadratic becomes $2(x^2 - \frac{3}{2}x + 5)$. This reduces the situation to dealing with a *monic* quadratic, which is a quadratic with leading coefficient equal to 1. So, let's worry about $x^2 - \frac{3}{2}x + 5$. The main technique now is to take the coefficient of x , which in our example is $-\frac{3}{2}$, divide it by 2 to get $-\frac{3}{4}$, and square it. We get $\frac{9}{16}$. We wish that the constant term were $\frac{9}{16}$ instead of 5, so let's do some

mental gymnastics:

$$x^2 - \frac{3}{2}x + 5 = x^2 - \frac{3}{2}x + \frac{9}{16} + 5 - \frac{9}{16}.$$

Why on earth would we want to add and subtract $\frac{9}{16}$? Because the first three terms combine to form $(x - \frac{3}{4})^2$. So, we have

$$x^2 - \frac{3}{2}x + 5 = \left(x^2 - \frac{3}{2}x + \frac{9}{16}\right) + 5 - \frac{9}{16} = \left(x - \frac{3}{4}\right)^2 + 5 - \frac{9}{16}.$$

Now we just have to work out the last little bit, which is just arithmetic: $5 - \frac{9}{16} = \frac{71}{16}$. Putting it all together, and restoring the factor of 2, we have


$$\begin{aligned} 2x^2 - 3x + 10 &= 2\left(x^2 - \frac{3}{2}x + 5\right) = 2\left(\left(x - \frac{3}{4}\right)^2 + \frac{71}{16}\right) \\ &= 2\left(x - \frac{3}{4}\right)^2 + \frac{71}{8}. \end{aligned}$$

It turns out that this is a much nicer form to deal with in a number of situations. Make sure you know how to complete the square, since we'll be using this technique a lot in Chapters 18 and 19.

2. Rational functions: these are functions of the form

$$\frac{p(x)}{q(x)},$$


where p and q are polynomials. Rational functions will pop up in many different contexts, and the graphs can look very different depending on the polynomials p and q . The simplest examples of rational functions are polynomials themselves, which arise when $q(x)$ is the constant polynomial 1. The next simplest examples are the functions $1/x^n$, where n is a positive integer. Let's look at some of the graphs of these functions:




The odd powers look similar to each other, and the even powers look similar to each other. It's worth knowing what these graphs look like.

22 • Functions, Graphs, and Lines

3. Exponentials and logarithms: you need to know what graphs of exponentials look like. For example, here is $y = 2^x$:




The graph of $y = b^x$ for any other base $b > 1$ looks similar to this. Things to notice are that the domain is the whole real line, the y -intercept is 1, the range is $(0, \infty)$, and there is a horizontal asymptote on the left at $y = 0$. In particular, the curve $y = b^x$ does **not**, I repeat, not touch the x -axis, no matter what it looks like on your graphing calculator! (We'll be looking at asymptotes again in Chapter 3.) The graph of $y = 2^{-x}$ is just the reflection of $y = 2^x$ in the y -axis:



How about when the base is less than 1? For example, consider the graph of $y = (\frac{1}{2})^x$. Notice that $(\frac{1}{2})^x = 1/2^x = 2^{-x}$, so the above graph of $y = 2^{-x}$ is also the graph of $y = (\frac{1}{2})^x$, since 2^{-x} and $(\frac{1}{2})^x$ are equal for any x . The same sort of thing happens for $y = b^x$ for any $0 < b < 1$, not just $b = \frac{1}{2}$.

Now, notice that the graph of $y = 2^x$ satisfies the horizontal line test, so there is an inverse function. This is in fact the base 2 logarithm, which is written $y = \log_2(x)$. Using the line $y = x$ as a mirror, the graph of $y = \log_2(x)$ looks like this:



The domain is $(0, \infty)$; note that this backs up what I said earlier about not being able to take logarithms of a negative number or of 0. The range is all of $(-\infty, \infty)$, and there's a vertical asymptote at $x = 0$. The graphs of $\log_{10}(x)$, and indeed $\log_b(x)$ for any $b > 1$, are very similar to this one. The log function is very important in calculus, so you should really know how to draw the above graph. We'll look at other properties of logarithms in Chapter 9.

4. Trig functions: these are so important that the entire next chapter is devoted to them.

5. Functions involving absolute values: let's take a close look at the *absolute value function* f given by $f(x) = |x|$. Here's the definition of $|x|$:


$$|x| = \begin{cases} x & \text{if } x \geq 0, \\ -x & \text{if } x < 0. \end{cases}$$

Another way of looking at $|x|$ is that it is the distance between x and 0 on the number line. More generally, you should learn this nice fact:


$$|x - y| \text{ is the distance between } x \text{ and } y \text{ on the number line.}$$



For example, suppose that you need to identify the region $|x - 1| \leq 3$ on the number line. You can interpret the inequality as “the distance between x and 1 is less than or equal to 3.” That is, we are looking for all the points that are no more than 3 units away from the number 1. So, let's take a number line and mark in the number 1 as follows:



The points which are no more than 3 units away extend to -2 on the left and 4 on the right, so the region we want looks like this:




So, the region $|x - 1| \leq 3$ can also be described as $[-2, 4]$.


24 • Functions, Graphs, and Lines

It's also true that $|x| = \sqrt{x^2}$. To check this, suppose that $x \geq 0$; then $\sqrt{x^2} = x$, no problem. If instead $x < 0$, then it can't be true that $\sqrt{x^2} = x$, since the left-hand side is positive but the right-hand side is negative. The correct equation is $\sqrt{x^2} = -x$; now the right-hand side is positive, since it's minus a negative number. If you look back at the definition of $|x|$, you'll see that we have just proved that $|x| = \sqrt{x^2}$. Even so, to deal with $|x|$, it's much better to use the piecewise definition than to write it as $\sqrt{x^2}$.

Finally, let's take a look at some graphs. If you know what the graph of a function looks like, you can get the graph of the absolute value of that function by reflecting everything below the x -axis up to above the x -axis, using the x -axis as your mirror. For example, here's the graph of $y = |x|$, which comes from reflecting the bottom portion of $y = x$ in the x -axis:



How about the graph of $y = |\log_2(x)|$? Using the reflection of the graph of $y = \log_2(x)$ above, this is what the absolute value version looks like:



Anyway, that's all I have to say about functions, apart from trig functions which are the subject of the next chapter. Hopefully you've seen a lot of the stuff in this chapter before. Most of the material in this chapter is used over and over again in calculus, so make sure you really get on top of it all as soon as you can!